Southeast Asian Illuminated Manuscripts: A Study of the Maritime Silk Road using Scientific Imaging & AI

Luke Butler

A thesis submitted in partial fulfilment of the
requirements of Nottingham Trent University
for the degree of Doctor of Philosophy.

March 2023

Copyright Statement

To My Grandmother, Sylvia, In Loving Memory

Abstract

In this project, the British Library's maritime Southeast Asian (SEA) illuminated manuscript collection was analysed via data visualisation and machine learning techniques to complement the holistic multimodal large-scale identification of artistic materials using visible to near-infrared (VNIR) reflectance spectral imaging (SI) and other complementary analysis techniques. Other techniques used included X-Ray fluorescence, external reflection-Fourier transform infrared, ultraviolet-visible-short wave infrared and Raman spectroscopies, all of which have shown to perform well for material identification when used together.

To condense large-scale spectral imaging data into manageable datasets, a new automated clustering and pigment identification informed grouping method was developed to reduce billions of VNIR spectra collected during large-scale SI surveys into smaller sets of unique spectral groups. Using this methodology meant that most artistic materials throughout the collection could be characterised and have their distributions mapped in less than 200 groups, where each group possessed a unique pigment mixture used within 18^{th}-19^{th} century maritime SEA. By using pigment mixture maps produced after grouping, 43 different unique pigment mixtures could be mapped and detected within 50 different manuscripts. One of the main findings showed an increasing European influence in the region as years progressed, but a tendency to mainly use traditional materials as late as the early-mid-19^{th} century, e.g. With a new detection of bone ash white being common in early Javanese manuscripts. The results offered by this study provided great insight into the use of artistic materials during this period and can be used now as a fundamental base for future research into maritime SEA and by extension the maritime silk road.

Additionally, a novel holistic multimodal clustering technique was also developed which allowed for the automated clustering of SI data created via different complementary techniques. And in addition to this, a new classification method was proposed which would allow for future large-scale SI data analysis and material identification to be performed as part of a hybrid clustering and classification approach.

Acknowledgements

Firstly, I would like to express my most sincere gratitude to my first supervisor Prof. Haida Liang, for her fantastic research opportunities, constant guidance, wisdom, and mentoring throughout the course of this PhD journey. I would additionally like to thank Dr Natasha Hodgson for her early insights into the historical applications of the research and her continued commitment as my second supervisor.

I would like to especially thank the British Library's Lead Curator for Southeast Asia, Dr Annabel Teh Gallop, for her own informative research insights and efforts in the arrangements for the many field trips required to collect the data seen throughout this thesis.

Furthermore, I'd like to express my gratitude to Dr Sotiria Kogou for her own outstanding research insights and essential help in performing data collection at the British Library. And I would like to additionally thank Dr Sammy Cheung for his efforts during the preparations and logistics required to carry out different research trips to the British Library also.

More generally, I would like to thank the British Library as an institution for providing support and access to the manuscript collections needed for me to carry out my research. And I also wish to express my gratitude to the Getty Conservation Institute for their collection of XRF mapping data used within this thesis.

I would also like to express my thanks to the entire current and previous ISAAC research group team members with whom I have had much support and many constructive conversations both about research, life, the universe and everything.

I wish to express my gratitude to my partner, my friends and my family for their support and patience throughout the writing of my thesis. Especially my mother, who has provided her unwavering support in the last year and during the COVID-19 pandemic.

And finally, I would like to express my gratitude to the Midlands3Cities (now M4C) doctoral training partnership and the AHRC which funded the PhD project, as without this I would never have had the opportunities I have today.

List of Publications

[1] Luke Butler, Sotiria Kogou, Alexander Hogg, Yu Li, Alessandra Vichi, Annabel Gallop, and Haida Liang "Machine learning analysis of complementary multimodal spectral imaging data from a large manuscript collection", Proc. SPIE 11784, Optics for Arts, Architecture, and Archaeology VIII, 1178405 (20 June 2020); https://doi.org/10.1117/12.2593918

[2] (eds) (31 Aug. 2023). Maps and Colours, Leiden, The Netherlands: Brill. Available From: Brill https://doi.org/10.1163/9789004467361 (IN PRESS)

Table of Contents

Chapter 1
Introduction

1.1 Background & Motivations

Maritime Southeast Asia is a region split between many islands, cultures, and religions, and has seen many interactions both within the archipelago and beyond throughout its history, each of which has led to many evolutions within the region towards how it exists today. While maritime Southeast Asia has its own long-standing and rich traditions throughout many different islands and cultures, many outside influences have contributed to the development of the region, including those from India, China, Europe, and the Islamic world (Tarling, 1992, 2008), which had such influence that Islam remains the primary religion linking most of the archipelago even to this day. With the many cultural influences brought by these external civilisations came further developments in many writing traditions, resulting in the adoption and adaptation of manuscript styles including manuscript illumination. Most surviving illuminated manuscripts found in the modern day from maritime Southeast Asia can be dated to the 17th-19th centuries, a time which saw many developments in the uses of artistic materials in different parts of the world. Many of the examples found today are finely illuminated, with many different artistic schools with different histories and cultures creating numerous manuscripts in many styles. However, little is known of the pigments and other artistic materials implemented during this period within the region and by extension, the use and trade of artistic materials between different regions within the Malay-Indonesian archipelago is also enigmatic. Unfortunately, by the late-19th century, many illuminated manuscript traditions had begun to fall into decline in the region, primarily due to the increasing popularity of printing in an industrialising world (A. Kumar & McGlynn, 1996). This means that today, the number of well-illuminated maritime Southeast Asian manuscripts is limited, and as such, very few studies have been performed to understand the types of artistic materials used in their preparation and production.

In the last hundred years or so, numerous organisations within Great Britain have sought to collect many of these manuscripts, and at the time of writing this thesis, the British Library currently holds one such collection containing a sizeable number of well-illuminated manuscripts. The collection, mainly derived from two major historic collections originally compiled by the British Museum and India Office Library, includes many different items from numerous islands within 18th-19th century maritime Southeast Asia. Such a collection is uncommon amongst many museums and libraries and therefore offers a unique opportunity to gain a greater understanding of the artistic practises used in maritime Southeast Asia, however, due to there being very little study into the materials used across the collection and the number of well-studied illuminated

manuscripts being limited over most of the world, this is a task which cannot be performed through observational and comparative studies alone. This, therefore, indicates a profound need for scientific methods to be employed which have shown in many past studies to allow for the identification of artistic materials, in turn providing the insight needed to draw conclusions about the evolution of trade, cultural exchanges and use of artistic materials between regions and along maritime trade routes over time.

Scientific Study of Artistic Materials

The technical study and scientific analysis of artistic materials is well-established and commonplace within the cultural heritage research community. In fact, it is now commonly accepted that scientific methods are crucial to the analysis, characterisation, and conservation of artworks. This is because using scientific techniques to perform material identification on manuscripts is the only reliable way to acquire information regarding the inks, paper, pigments, binders, and other materials, which are useful for giving insights into the artistic practises of different cultures and regions throughout history. As the materials used in the preparation and production of maritime Southeast Asian illuminated manuscripts have been studied very rarely, the scientific examination of the British Library's collection, which has only been performed once before by Burgio et al. (1999), offers the ability to build a fundamental basis for cultural and historical research into the use of pigments and other materials within the archipelago throughout the 18^{th}-19^{th} centuries. Therefore, throughout this thesis different scientific data collection and analytical techniques are implemented to study the use of pigments and other materials in numerous maritime Southeast Asian manuscripts.

To acquire the data necessary for informing the use of artistic materials, the normal practise is to perform a holistic multimodal study, where many different techniques are employed together towards a combined goal. This is typically necessary because there is no one technique that can characterise every property required to identify materials, which implies that techniques must be used which best complement each other. In modern studies, many different techniques such as fibre optic reflectance spectroscopy, X-Ray fluorescence, Raman spectroscopy, and others, have all been implemented in the pursuit of performing material analysis on cultural heritage objects. As technology has progressed in more recent decades, most techniques have now developed to be able to collect spectral image-based scientific data non-invasively, meaning that millions of datapoints useful for characterising materials can be recorded at the same time while avoiding any potential damage to objects. Furthermore, modern systems can also perform data collection in-situ, meaning that for large-scale scientific studies of manuscripts, different instruments can be brought to collections instead of having to transfer potentially fragile items across cities, countries or continents. These factors not only make the analysis of a collection safer, but also mean that a greater number of items can be analysed more efficiently as there are very few logistical problems that may occur on account of moving many manuscripts to perform large-scale analysis.

Artificial Intelligence and Machine Learning Necessity

When performing heritage studies covering vast geographical regions, large volumes of data are required to make statistically reliable conclusions about the use of artistic materials. This way if any trends in the implementation of materials or styles can be found in certain regions, it can be assumed that these artistic practises may be the norm. However, performing scientific data collection at such a large scale to acquire this information introduces new problems where data becomes too large or complex to analyse efficiently. Therefore, the need for computational techniques that can aid in improving the efficiency of analysis become increasingly necessary. In recent years, machine learning technologies have demonstrated their capability to aid in this pursuit within cultural heritage research, and for the problems encountered within this thesis, they can offer unique ways to allow for the analysis of 'big data', which will only grow to become even bigger as more data is collected over time and in greater abundance. Machine learning covers a swath of computational techniques that focus on performing data analysis by imitating how we, as humans, learn, and is largely considered to fall under the umbrella of artificial intelligence (AI). While machine learning was originally conceived in the 1940s (McCulloch & Pitts, 1943), advances in memory storage and processing power over the last few decades have made machine learning widely available, and as more powerful computers are now more accessible than ever, machine learning can be seen used in many aspects of modern life, anywhere from Google's search engines to Tesla's self-driving cars. In fact, in 2021 a McKinsey global survey which monitored the working practices in 1,843 different companies worldwide found that 56% of respondents had adopted and were using AI in at least one function (McKinsey Global Institute, 2021). With machine learning becoming more popular in various industries and applications, there are now many processes which can be implemented into different scenarios to solve large and complex data analysis problems, including dimensionality reduction for computer vision and visualisation, reinforcement learning for robotics, or predictive modelling for stock market forecasting. However, for the task of performing the large-scale scientific data analysis required for this thesis, the most useful machine learning-based techniques to use are unsupervised clustering and supervised classification, both of which can be implemented to reduce the size of large datasets and or perform automated material analysis. As such, the use of both techniques will be presented throughout this thesis in pursuit of performing large-scale data analysis for manuscript collections and paper-based paintings (with both being discussed further in sections 1.3 and 1.4 of this chapter). In most chapters of this thesis, however, the focus is on the use of clustering applied to spectral imaging data, due to its capabilities of simplifying complex datasets into more easily interpretable information.

1.2 Holistic Scientific Study of Illuminated Manuscripts

The study of artistic materials used within illuminated manuscripts or artworks normally requires the collection of different types of data so that conclusive material analysis can be performed. This approach is normally taken as the implementation of different complementary techniques can be proven to not only provide the data necessary for the identification of materials (Delaney et al., 2014; Kogou et al., 2015; Miliani et al., 2010), but they can also be used to guide the interpretation of results as well, as one technique can provide information to inform others, improving the overall accuracy of the study (Kogou, 2017). For the studies of manuscripts and other paper-based paintings performed throughout this thesis, there are several non-invasive and non-destructive techniques which can be used to provide the information to determine artistic materials, therefore the analysis techniques used in this thesis are discussed in the following sections.

1.2.1 Visible-Near Infrared (VNIR) Reflectance Spectroscopy

One of the most widely used and accessible techniques for analysing artistic materials is visible-to-near-infrared (VNIR) reflectance spectroscopy. When visible incident light meets a solid object, it will normally be split into different components, where it can either be absorbed or reflected, transmitted, or scattered (see Fig. 1.1).

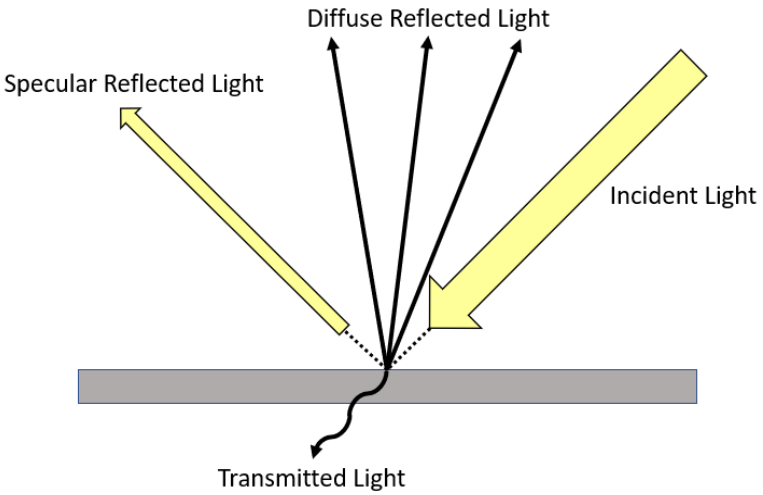

Fig.1.1 Basic diagram illustrating the reflection and transmission of incident light in VNIR reflectance spectroscopy.

The proportion of light belonging to any one of these components in the VNIR range, which covers roughly 400-1000nm in wavelength, is based on many different factors including colour, transparency, particle size, molecular structure, or surface topography. By monitoring the amount of diffuse light reflected from an object, properties unique to specific materials, or mixtures of materials, can be inferred and compared against references to identify pigments, inks, dyes, etc.

These measurements of the reflected light from a material are known as reflectance, and are wavelength dependent, implying a variation in responses for different unique materials at different wavelengths of light. In general, VNIR reflectance is calculated by comparing the reflected light from an object against a reference target with a known response at different wavelengths, thereby producing a relative reflectance $R(\lambda)$ with values between 0 and 1, which can be used for analysis as different materials will reflect light in characteristic ways. VNIR reflectance has proven to be a powerful tool and has shown to be pivotal in the identification of many different materials in a large variety of studies covering the analysis of statues, paintings, and illuminated manuscripts, and has been used in many other aspects of art history and conservation research such as monitoring laser cleaning, bronze disease, and more. Typically, the technique commonly utilised in heritage studies is fibre optic reflectance spectroscopy, but there are other methods which can be implemented.

Fibre Optic Reflectance Spectroscopy (FORS)

Fibre Optics reflectance spectroscopy (FORS) is a powerful tool used for material analysis in cultural heritage with an extensive history of use for multiple decades in the study and examination of artworks (Leona & Winter, 2001). It utilises basic equipment and works from a simple concept: to direct a source of light onto a small point using a fibre optic probe and monitor the reflected light using a spectrometer to produce a reflectance spectrum. A basic diagram illustrating the setup of a FORS system can be seen in Fig. 1.2.

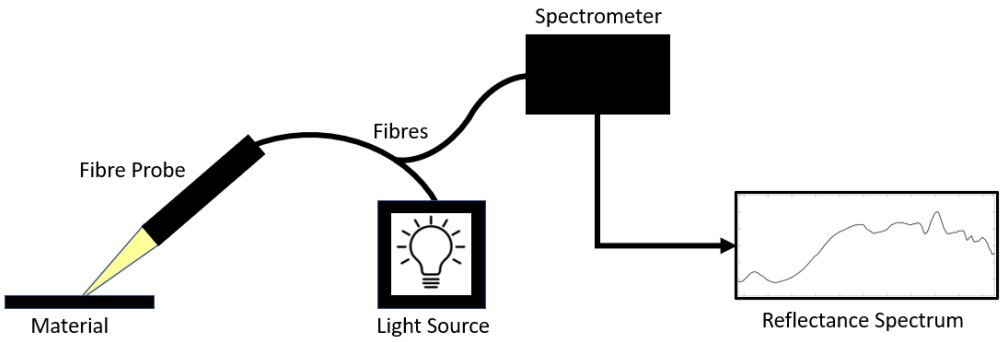

Fig.1.2 Diagram showing the basic setup for performing fibre optic reflectance spectroscopy.

FORS has demonstrated the ability to identify inorganic and organic pigments, dyes, and mixtures making it a great candidate for the analysis of manuscripts as it can provide reflectance spectra which can detect fine absorption features, such as those seen in red dyes. FORS has already seen use in multiple studies of pigments and illuminated manuscripts and is therefore clearly useful for the identification of materials in this study, it can however be a technique that is not efficient for

data collection over large collections, which is one of the main goals of this thesis. Therefore, spectral imaging-based alternatives were used to gather large amounts of data first, where FORS was then implemented strategically to identify specific materials where required.

1.2.2 VNIR Reflectance Spectral Imaging

Spectroscopy can provide useful information in the study of artworks and manuscripts, but when considering the examination of complex large-scale datasets, both data collection and analysis may become a colossal task to perform if research questions require spectroscopy point measurements to be performed over thousands of different locations, therefore, alternative time and work efficient methods are desirable. In more recent years, developments in optical and other scanning techniques have led to the more common use of spectral imaging systems which can collect large amounts of data efficiently (Liang, 2012). Frequently used in heritage science, spectral imaging provides an increased volume of data by combining imaging with spectroscopy, allowing for the creation of three-dimensional volumetric data cubes (or image cubes), where every pixel within the two-dimensional spatial profile of an object also contains a spectrum unique to the materials present at the pixel location. Not only does this allow for the collection of thousands, or millions of spectra simultaneously, but it also provides information that can eventually allow for the large-scale identification and mapping of artistic materials on an artwork or a manuscript. Furthermore, with the correct processing in the VNIR, the true colour information of an object can also be produced by using this reflectance information, allowing for spectral imaging to provide colour-accurate digital images in addition to scientific data analysis. When performing spectral imaging in cultural heritage, the properties of collected spectra are heavily affected and defined by the instruments used, and in VNIR spectral imaging, many different non-invasive approaches can be implemented. These include but are not limited to techniques which use filtered illumination, certain instrumental optical filters, and specific gratings (Cucci et al., 2016; Liggins et al., 2022) which can exploit the diffraction of broadband light into its different constituent wavelengths. In this thesis, spectral imaging systems using filtered illumination and optical filters are both used to collect data for different paintings and are therefore the only techniques discussed in greater detail in this section.

Wavelength Scanning - Filtered Illumination

One of the most basic and convenient forms of performing spectral imaging in the VNIR regime is to synchronise image capture with controlled illumination for several wavelengths of light. In some disciplines, filtered illumination techniques are normally used to visualise the individual spectral bands one at a time as each can provide clues as to what materials have been used on a manuscript, or can reveal objects hidden to the eye in ordinary lighting conditions. For example, recording a single image when illuminating a manuscript page with UV light can allow for the

capture of UV fluorescence from inks, where fluorescence can indicate the use of materials like iron gall. Additionally, using individual infrared bands allows for penetration through layers of paint or dirt, providing images which can often show clear underdrawings or other substrate features (Duffy, 2018; McGillivray & Duffy, 2017). The difficulty with these studies however is that the techniques do not always offer the ability to acquire spectra images, meaning that the identification of other materials is not possible. However, while many studies will only use scientific imaging in this way, with ideal controlled conditions during data collection followed by applying appropriate processing techniques, the acquisition of accurate reflectance spectral information can still be performed and therefore spectral imaging data cubes can be produced for use in material identification.

Wavelength Scanning – Filtered Optics System

Reflectance spectral imaging data can also be acquired by placing different optical filters into a system so that only certain frequencies of light are permitted to travel through an optical system and into a detector. By scanning the same object through multiple different filters, different image channels can be generated representing the reflectance for different wavelengths of light. When combined, the channels can be sorted into a spectral imaging data cube. There are multiple ways in which this can be performed, but the most common approaches normally involve a mechanical way of swapping between filters. For example, in Liang et al., (2014), the same spectral imaging system used throughout this thesis (PRISMS) was implemented, where data is recorded by placing an electronically controlled filter wheel between the object of interest and the detector optics so that ten filters with regular bandwidths can be used to generate spectral imaging data in the VNIR regime under controlled illumination. Filtered optics systems are normally of a lower spectral resolution than some other techniques (e.g. grating-based line scanning approaches) commonly referred to as "hyperspectral", however, they have still seen much use in cultural heritage studies and are well documented to provide reflectance spectra which can be used to identify common pigments in paintings, as most pigments have broad spectral features, making the technique an excellent candidate for the analysis of manuscripts. The only situations in which higher spectral resolution would be required are to identify materials with fine absorption features, for which the most implemented approach taken in this thesis was to perform point measurements using FORS in specific areas of interest after spectral imaging was performed, as the same information could be acquired for pigment identification without having to implement higher spectral resolution imaging.

Colour Image Production using Reflectance Spectral Data

Spectral imaging in the VNIR range can also be used to produce colour accurate images of objects under different lighting conditions by deriving a set of device-independent tri-stimulus values

from the spectral reflectance (McLaren & Rigg, 1976). This can be achieved by first converting to XYZ colour space through the following equations:

$$X = \frac{k}{N}\int_0^\infty R(\lambda)I(\lambda)\bar{x}(\lambda)d\lambda$$
$$Y = \frac{k}{N}\int_0^\infty R(\lambda)I(\lambda)\bar{y}(\lambda)d\lambda$$
$$Z = \frac{k}{N}\int_0^\infty R(\lambda)I(\lambda)\bar{z}(\lambda)d\lambda$$
$$\frac{k}{N} = \frac{100}{\int_0^\infty I(\lambda)\bar{y}(\lambda)d\lambda}$$

Where R is the spectral reflectance of any material recorded by the system, I is the emission spectrum of the illumination, k is the scaling factor, and $\bar{x}, \bar{y}, \bar{z}$ are the colour-matching functions describing the response of the CIE 1931 2° standard observer (Cie, 2004). Once calculated, XYZ tri-stimulus values are typically then converted into CIELAB (McLaren & Rigg, 1976) colour space as it enables the accurate measurement and comparison of colours where numerical differences roughly correspond to how humans perceive changes in colour. It does this by defining colour not by three channels of different colours, but instead by modelling its channels around the unique perception of colours in human vision, which consists of measurements for the perceptual lightness or Luminance (L^*), red/green response (a^*), and blue/yellow response (b^*). The conversion from XYZ to CIELAB (or $L^*a^*b^*$) can be achieved by using the following:

$$L^* = 116f\left(\frac{Y}{Y_n}\right) - 16$$
$$a^* = 500\left(f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right)$$
$$b^* = 200\left(f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right)$$
$$f(x) = \begin{cases} \sqrt[3]{t}, & if\ t > \delta^3 \\ \frac{t}{3\delta^2} + \frac{4}{29}, & otherwise \end{cases}$$

Where $\delta = \frac{6}{29}$, X, Y, Z are the tri-stimulus values of a material and X_n, Y_n, Z_n are the values for a white reference illuminant, normally the standard D65 illuminant. This conversion will work for

a single reflectance spectrum, however when applied to spectral data cubes can be used to generate entire colour images. From CIELAB colour space, different other colour spaces can be achieved, such as the more common sRGB (IEC, 1999), with colours able to be simulated for different lighting conditions if the emission of the light source during data capture of spectral data is known. Throughout this thesis, colour is used in different ways to aid in the analysis of different paintings and is central to a novel machine learning-based clustering method developed for the analysis of the maritime Southeast Asian manuscript collection, where CIELAB is used directly with spectral reflectance.

1.2.3 X-Ray Fluorescence (XRF) Spectroscopy

Another exploited phenomenon used within cultural heritage is X-Ray Fluorescence (XRF). XRF is a non-invasive technique commonly used to determine the elemental composition of different materials through the identification of characteristic 'peaks', guassian features in an energy spectrum which correspond to different elements. In heritage studies, many different pigments, inks, substrates, or other materials often possess characteristic elemental composition and therefore can be identified using XRF. It functions by directing X-rays (known as primary X-rays) to an object and reading back fluorescent X-rays (secondary X-rays) caused by the incident primary X-ray's interaction with matter at the atomic level.

More specifically, it occurs when the electrons in different orbitals of an atom transition from higher to lower energy states, usually caused by electron bombardment from the primary X-rays or photon absorption from gamma rays. The basic premise is that if an electron is ejected from the innermost (K) orbital to an outer shell, the atom enters a higher energy excited state, leaving an electron vacancy in the K shell which needs to be filled to return the atom to its normal state. To return the ionised atom to its stable state, another electron from one of the outer shells must transition to fill this K shell, where, for example, in a K-L (Kα) transition, this involves an electron moving from the L to K shells. When the higher energy electron cascades down to a lower shell, energy equivalent to the binding energy required to move from K to L is released from the atom in the form of a secondary fluorescent X-ray (Kα line). The process of filling an electron vacancy in the K shell in turn creates more vacancies in higher orbitals, which also require filling to maintain the stability of the atom. The full process, therefore, results in many transitions which together release secondary x-rays with energies characteristic of specific chemical elements, allowing for the identification of elements in a material. In some cases, K-L transitions are not the only permitted transitions. For example, for atoms with large numbers of electrons, the K shell vacancy may instead be filled by an electron transitioning from the M orbital, a K-M transition, which causes the emission of the Kβ line. Many transitions are also permitted including those from the L, M, N, shells etc. Some of the common transitions seen in XRF analysis with their

associated Siegbahn and IUPAC notation is included in Fig. 1.3, which describe specific energy transitions for different chemical elements.

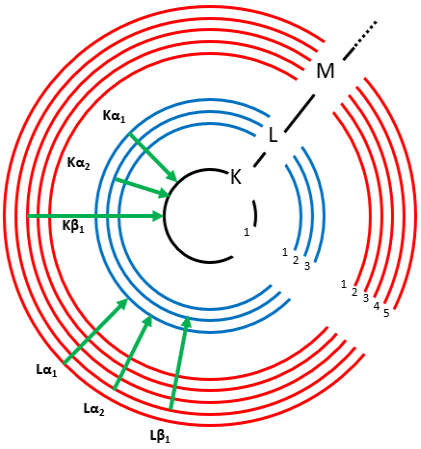

Some Main Characteristic X-Ray Transitions and Notation			
Transition Shells		Siegbahn Notation	IUPAC Notation
Source	Destination		
L-3	K	Kα - 1	$K\text{-}L_3$
L-2	K	Kα - 2	$K\text{-}L_2$
M-3	K	Kβ - 1	$K\text{-}M_3$
M-5	L-3	Lα - 1	$L_3\text{-}M_5$
M-4	L-3	Lα - 2	$L_3\text{-}M_4$
M-4	L-2	Lβ - 1	$L_2\text{-}M_4$
M-5	L-3	Lβ - 2	$L_3\text{-}N_5$
N-4	L-2	Lγ - 1	$L_2\text{-}N_4$
N-7	M-5	Mα - 1	$M_5\text{-}N_7$

Fig.1.3 Left: Diagram showing typical K, L, M shells with some transitions included as arrows. Right: A table covering some of the main characteristic X-ray transitions and their notation.

XRF Energy Range Detection

In standard use of most portable XRF systems (seen implemented in this thesis) the detectable range for chemical elements normally spreads from Z=14 (Silicon) to Z=92 (Uranium). As the materials in cultural heritage are not likely to lay above Uranium in the period table of elements, the upper limit on the detectable range is not a huge problem. However, there are many materials in cultural heritage which may contain elements with Z<14, which would not be identifiable due to the absorption of characteristic x-rays by the air at lower energy ranges. To address this, most modern XRF analysis instruments can perform a Helium purge, where the air around a material being analysed is replaced with Helium gas, permitting the detection of elements down to Z=11 (sodium). Even with the implementation of a Helium purge, however, there are still many lighter elements that are typically not detectable using portable XRF systems. These include elements which typically correspond to organic pigments such as Carbon, Nitrogen and Oxygen.

Bremsstrahlung/Background Continuum

Performing XRF measurements always results in the creation of bremsstrahlung radiation from either the XRF system or the surrounding environment. Bremsstrahlung primarily presents itself as the background continuum in an XRF energy spectrum, and its profile can change between systems and with the voltage supplied to the XRF ray tube and therefore the energy of the emitted X-Rays. For most of the thesis the bremsstrahlung continuum had minimal effect on the identification of elements, however, there were some scenarios where the signal representing certain elements was weak enough that the background can be seen to dominate most of an XRF

spectrum, requiring continuum subtraction algorithms to be applied to the data prior to analysis, as demonstrated in chapter 4.

Rayleigh and Compton Scattering

During the process of performing XRF, Rayleigh and Compton scattering can occur as a side effect of photons interacting with matter and are usually presented as a set of peaks in the XRF spectrum. In Rayleigh scattering, this involves the elastic scattering of a primary photon, meaning that it does not lose energy and therefore can be detected as a set of peaks with energies equal to the material used in an X-ray tube (typically rhodium, tungsten, or silver). In Compton scattering, inelastic scattering occurs within the sample being analysed, meaning that a primary photon released from an XRF system will interact with matter, resulting in scattering where the photon returns to the detector with a lower energy. Therefore, the result of Compton scattering is that a lower energy photon than the incident X-Ray beam will be detected therefore visible on any XRF spectrum as peaks which are positioned at slightly lower energies (keV) than those for Rayleigh Peaks.

Sum Peaks

Normally occurring in scenarios where a single element dominates the entire XRF spectrum, it is possible for multiple secondary photons to enter an XRF detector simultaneously, producing a signal with energy equivalent to the sum of the different secondary X-rays. Due to the higher occurrence of their associated electron transitions, sum peaks most often occur because of Kα/Kα, Kα/Kβ, Kβ/Kβ pairs entering the detector at the same time, but any combination of simultaneous readings can happen. If this phenomenon happens repeatedly over a full XRF measurement, a new, but often small, "sum peak" may be detected which can cause problems when attempting to identify characteristic XRF lines of certain trace elements.

Escape Peaks

In many circumstances, very strong energy peaks can be detected in an XRF spectrum, resulting in the formation of other artefacts such as escape peaks. Escape peaks are produced because of XRF fluorescence photons passing through the detector crystal with significant enough energy to cause further x-ray emission. Released X-rays are normally reabsorbed by the detector and a spectrum produced as normal; however, it is also possible for these photons to escape the crystal. As most modern detectors are made from silicon, the x-rays normally released during this process are Si-Kα. So when this occurs in modern systems, the energy of the photons collected by the detector is equivalent to the energy of the original fluorescent photon minus the energy of the Si-Kα photon which has escaped from the system, where total energy collected: $E_{Detected} = E_{Initial} - E_{Si-K\alpha}$. In an XRF spectrum, this means that we would expect to see two distinct peaks approximately 1.74keV apart (as $E_{Si-K\alpha}$≈1.74keV) for chemical elements that have extremely high counts in the spectrum.

Other Instrumental Peaks

In addition to those previously mentioned, it is common for an XRF spectrum to possess peaks for chemical elements which can be found in the air or X-ray tube used to produce X-Rays. For example, a common material to use for the anode is Rhodium (Rh), which means that if XRF spectra are produced, there will always be a Rhodium peak present at both the lower and higher energy levels. Likewise, in air, Argon (Ar) is the third most abundant element, so when performing XRF it is not uncommon for Argon to exist as a large peak in the XRF energy spectrum (which is why helium purge is often used in environments where the Argon peak may cause difficulties in analysing materials).

1.2.4 XRF Mapping

XRF mapping has become more common in the last few decades due to the improvements in computational power and increased number of systems available commercially. XRF Mapping is typically performed by performing variations of single-point scanning but also be performed with on-the-fly techniques (Alberti et al., 2017) which can result in much faster data acquisition. Regardless, most mobile systems used in studies such as those seen in this thesis collect data by moving the head of the XRF system (containing the X-Ray tube) to cover an area over a painting or manuscript, therefore, allowing for convenient non-invasive scanning without having to move the object. XRF mapping can result in data which is similar in structure to any other form of spectral imaging, where the 2D image plane is made up of individual pixels, with each containing its own XRF spectrum. The main difference however is that with other techniques, units of frequency or wavelength are used to cover continuous features over many different regimes of the electromagnetic spectrum, but with XRF, sharp peak-like structures covering different energies will be observed. Because of this, many applications using XRF can be satisfied by extracting the pixels at specific peaks, to generate elemental composition maps and illustrate where certain materials may exist with the use of different post-processing and data visualisation techniques. In cultural heritage studies, these elemental composition maps can provide important qualitative information about the spread of artistic materials used over an artwork. Such maps can illustrate where there may be different inorganic pigments or different layers of paints on a painting, providing not only material identification but also insight into the artistic methods used by artists in the past (Mazzinghi et al., 2021; Romano et al., 2017). Furthermore, as X-Rays can penetrate to much greater depths than some reflectance-based techniques commonly used in cultural heritage, XRF can provide information for materials present underneath the surface.

1.2.5 Fourier Transform Infrared (FTIR) Spectroscopy

Like VNIR reflectance spectroscopy, FTIR spectroscopy can also be used to perform the chemical characterisation of materials, making it a valuable tool for cultural heritage studies (Galeotti et al., 2009; G. L. Liu & Kazarian, 2022). FTIR functions on the principle that most vibrational interactions within molecules occur with fundamental vibration frequencies characteristic of specific chemical bonds in organic or inorganic materials. Many molecular vibrational interactions most commonly occur in the mid-infrared (MIR) region of the electromagnetic spectrum, which expands from 4000-500cm^{-1}, but the full range of a typical FTIR spectrum can often also include the near-infrared depending on the spectrometer/system being used. FTIR spectroscopy typically involves using an interferometer to acquire the infrared spectrum of a material. This is typically performed using setups such as the Michelson interferometer (see Fig. 1.4), whereby an interferogram is produced from interference between reflected infrared light after interacting with a material. The interferogram is further subjected to a Fourier Transform which can produce the transmission or absorption spectrum of a material if the infrared light has interacted with an object. To identify materials in use in cultural heritage, FTIR results can be compared with databases of known materials, allowing FTIR to perform material identification if produced spectra possess matching features to standard references.

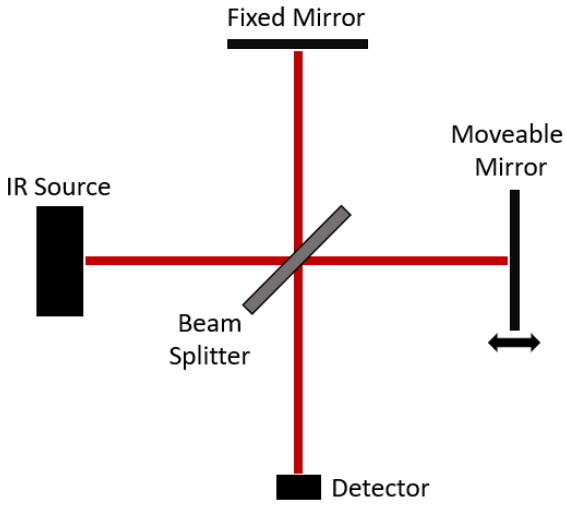

Fig.1.4 Basic example of a Michelson interferometer commonly used to perform FTIR spectroscopy.

There are however still limitations of the technique which mean that FTIR must be used in conjunction with other methods. For example, in many scenarios FTIR spectra can be difficult to analyse for complex chemical matrices, where multiple absorption bands caused by constituent compounds within a material may overlap one another, resulting in indiscernible spectral features which cannot be properly compared with references to identify materials. Furthermore, certain materials will also be indistinguishable using FTIR alone if they possess similar chemical

structures, as this will result in spectra with absorption bands in the same positions. FTIR is normally performed using either transmission or reflection. In many cases however, such as in the non-destructive analysis of manuscripts, using FTIR in its 'reflection mode' is the only way that analysis be performed in a completely non-destructive way, in contrast to transmission which typically requires a small, microscopic sample to be taken from an object. FTIR in its 'reflection mode' can be performed in different ways, either using Diffuse Reflection (DRIFT), Attenuated Total Reflection (FITR-ATR), or Specular Reflection (also referred to as ER-FTIR), each of which can be implemented in different applications/scenarios and offer different advantages and disadvantages depending on their usage. Throughout this study, FTIR is only used in its external reflection mode, where it was positioned over numerous areas of manuscript illumination to aid in the identification of pigments.

External Reflection FTIR (ER-FTIR)

ER-FTIR is a fully non-invasive method which gathers spectral information through the monitoring of specular reflection. It does not require sampling or any direct contact with materials, as is the case with other techniques such as FTIR-ATR, making it a powerful tool for the in-situ analysis of illuminated manuscripts. However, even though ER-FTIR is the most useful technique for analysing materials non-invasively in cultural heritage, it is not without its own difficulties. ER-FTIR-produced spectra are well known to suffer from increased complexity due to band distortions, in the form of S-shaped distortions, band inversion and changes in intensity ratios, leading to increased intensity of weaker bands (such as combination bands or overtones). These changes can occur due to differences in particle size, surface roughness, absorption, or refractive index of artistic materials and their substrates or binding media; each of these can affect the penetration of incident infrared light, which in turn alters spectral features. This means that spectra may appear differently in various parts of the infrared range in comparison to reference spectra, affected by unique conditions in which the material was prepared or applied. Consequently, ER-FTIR spectra must be carefully compared against material reference databases from other techniques such as FTIR-ATR, and in many cases may not have entirely matching spectra, making pigment identification a complex process. Regardless, ER-FTIR is the best FTIR-based candidate for performing material analysis of illuminated manuscripts as the positives of being able to perform data collection non-invasively far outweigh the difficulties of interpreting the collected spectra. As a result, throughout this thesis, ER-FTIR was implemented in many different circumstances to produce spectra for the purposes of pigment identification and can be seen in use in chapters 3 and 4.

1.2.6 ER-FTIR Mapping (FTIR Spectral Imaging)

In more recent years, FTIR mapping techniques have become more prevalent in material analysis studies and are proving to be even more useful than performing single measurements due to their ability to generate maps which can illustrate the distribution of different compounds on an object's surface. In many modern FTIR systems, ER-FTIR mapping can be used to non-invasively produce FTIR spectral images which share a structure like that of reflectance and XRF spectral imaging, meaning that a data cube would possess a two-dimensional array of pixels where every pixel possesses its own FTIR spectrum.

Typically, ER-FTIR mapping is performed via point scanning, where individual pixels are recorded one at a time and recombined after data collection to produce a spectral data cube. While this is a powerful technique for performing material analysis, the process of collecting an entire ER-FTIR map pixel by pixel can be time-consuming, implying that it would be unwise to attempt to use it for more than just small, controlled areas. In addition, the mapping can be sensitive to surface topography changes, as for ER-FTIR it is important that the materials being investigated are always positioned at the correct focal point from the system. If this is not the case certain spectra collected from different positions in a map may not be true to their actual location and may also result in the creation of artefacts during the recording of spectra, making the spectral comparison with material references difficult. In this thesis, ER-FTIR mapping is used in chapter 4 to perform focussed pigment identification with other techniques on a Chinese Export painting.

1.2.7 Raman Spectroscopy

Raman spectroscopy is another technique proven to be useful for identifying artistic materials in cultural heritage (Casadio et al., 2016; Galeotti et al., 2009; McCreery, 2000), and similarly to FTIR can provide information on the molecular vibration and chemical bonds of different materials. More specifically, Raman spectroscopy exploits the effects of inelastic scattering of a monochromatic light source (normally a laser in the visible, near-infrared, or near-ultraviolet range). When monochromatic light of a specific wavelength is incident on a material, it can sometimes be scattered at the same frequency. This scattering, known as Rayleigh scattering (briefly mentioned in 1.2.3 for XRF), occurs if both; the particle size within the material is smaller than the wavelength of incident light, and the electrons return to their ground state after temporarily transitioning to a virtual state. In addition to Rayleigh scattering, molecular vibrations can also occur due to the transfer of energy from incident photons. When both Rayleigh scattering and these molecular vibrations occur simultaneously, the two can result in different transitions which include the vibrational state. These different transitions are known as Raman scattering and take the form of two transitions, namely stokes and anti-stokes, where the energy is lower or

higher than Rayleigh Scattering depending on if the energy transition cascades to the vibrational state or ground state, a diagram illustrating these transitions can be seen in Fig. 1.5.

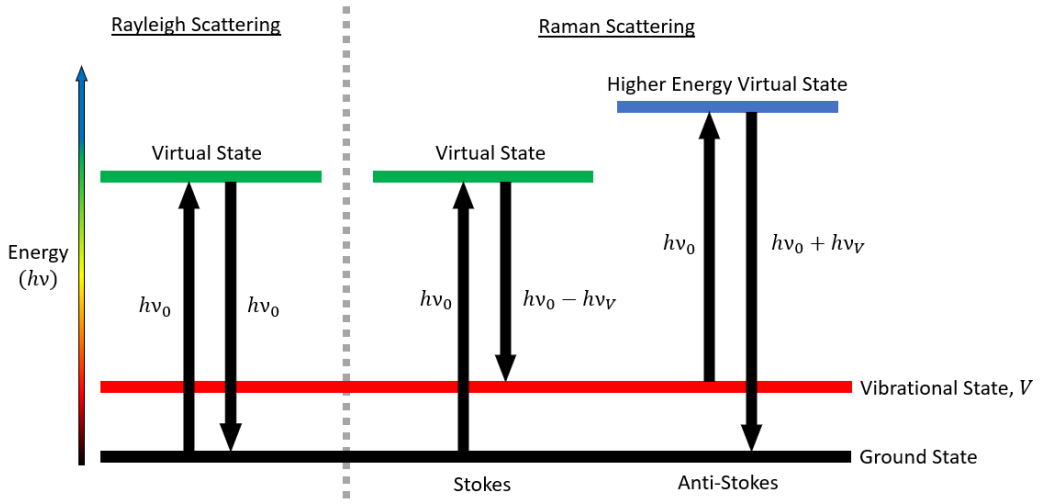

Fig.1.5 Diagram showing some of the energy transitions between different states for Rayleigh and Raman scattering. Both Stokes and Anti-stokes are included in this example.

Raman scattering can be monitored using a spectrometer to identify a Raman Spectrum, with emission occurring at different wavenumbers from the incident light, commonly referred to as Raman shift, where:

$$Raman\ Shift\ (cm^{-1}) = frequency_{incident}(cm^{-1}) - frequency_{raman\ scattered}(cm^{-1})$$

Emission at certain Raman shifts is characteristic of specific molecules and therefore can be used to identify materials by comparing the Raman spectra with reference databases. In cultural heritage, Raman spectroscopy has been performed in multiple ways, including micro-Raman (Caggiani & Colomban, 2018), and remote Raman (Li et al., 2019). In this thesis, Raman spectroscopy is used multiple times to aid in the identification of certain pigments and was essential in confirming the presence of certain pigments within different mixtures where constituent mixture components were ambiguous.

1.2.8 Kubelka-Munk (KM) Model for Pigment Identification

When attempting to identify paint mixtures during the analysis of VNIR reflectance spectra, it can be difficult to ascertain the constituent pigments used to create certain hues and shades. As such, different mathematical models can be used to estimate which materials may be used together to create certain mixtures. The technique used within this thesis is the Kubelka-Munk (KM) model (Kubelka, 1948; Kubelka & Munk, 1931), which can be used to model different mixtures by combining sets of VNIR reflectance reference spectra with varying concentrations. KM theory, in its most basic form in diffuse spectroscopy applications, can be used to monitor the reflectance and scattering of light in an infinitely thick coating by modelling the coating as an infinite number of infinitesimally small paint layers. In this scenario, KM theory can be described by a single equation:

$$\frac{K}{S} = \frac{(1 - R_\infty)^2}{2R_\infty}$$

Where R_∞ is the spectral reflectance for a paint layer with infinite thickness, K is the wavelength-dependent effective absorption, and S is the wavelength-dependent scattering coefficient. To identify manuscript pigments or other material properties in this thesis, KM theory can be used to determine potential mixtures by virtually mixing reference spectra and comparing the outcomes with VNIR spectra extracted directly from reflectance spectral imaging data or FORS. By changing the concentrations for known pure pigments and using them as free parameters in the KM model, simulated mixtures can be compared against unknown collected spectra and fitted using a non-negative least squared fitting algorithm (Liang, 2012). If a simulated spectrum has a good fit, it will have a high cross-correlation score with the unknown mixture as this will define closely matching spectral features.

With the combination of multiple complementary techniques in use with KM theory, it is likely that most mixtures could be identified within an illuminated manuscript. However, it is important to note that in some circumstances this may still not be evidence enough to perform confident pigment identification, as KM theory works off the assumption that the paint mixtures are highly scattering, homogenous and have particle sizes much smaller than the thickness of the paint layer, which is not always the case in practise. In many scenarios, variations can occur in paint layers which can affect the spectral reflectance and validity of applying KM, however, even in some scenarios where KM theory does not provide an exact representation of the mixture, it may still provide a "best fit" which may contain similar spectral features which can still allow for some identification when used with other techniques, so long as the analysed materials are not too highly absorbing or transparent.

1.3 Unsupervised Clustering Techniques in Machine Learning

To accelerate the process of data analysis after performing large-scale scanning of cultural heritage objects, unsupervised clustering methods can be implemented to automatically reduce billions of collected reflectance spectra into smaller numbers of distinct spectral groups. While there are numerous clustering techniques which can be used to assign data into highly controlled groups, many of the methods which may at first demonstrate good potential for analysing spectral data may in fact not be viable options for large-scale studies due to their heavy computational requirements. When performing clustering for larger datasets, it is important that both accuracy and time efficiency are well balanced, however many unsupervised techniques often require significant parameterisation and optimisation before clustering can be implemented, and even then, the algorithm may take days to complete. If a general clustering algorithm is to be designed to allow for the accurate and efficient grouping of reflectance spectra for entire illuminated manuscript collections, it would be wise to use a clustering approach that requires very few or no parameters and can perform clustering in a time-efficient manner. Because of this, in the following sections, only three clustering techniques used in various parts of the thesis are explored in detail, as these offered the ability to perform clustering for specific purposes both efficiently and accurately while requiring very little optimisation and normally only a single input parameter once the most optimal settings had been computed.

1.3.1 K-Means Clustering

Probably the most well-used procedure for performing unsupervised clustering is the k-means algorithm. K-means clustering can group data into a fixed number of clusters by identifying a certain number of cluster centroids in feature space, the N-dimensional space where datapoints exist, which best describe all unique groups of similar data points. K-means works by initially positioning the centroids at randomly selected data points and then iteratively updating the position of these centroids by monitoring and optimising the cluster sum of squares error (SSE) as new data points are assigned to clusters, where SSE is calculated by:

$$SSE = \sum_{i \in k} (x_i - \bar{x}_k)^2$$

Where x_i is a datapoint belonging to a cluster k and $\bar{x}$ is the mean value of the cluster calculated from all datapoints in k. Centroids are continuously updated until either a user-defined number of iterations is reached, or the positions of the cluster centroids relative to their associated data points no longer change, a point known as convergence. Multiple versions of k-means clustering have been designed since its original conception, and the three most widely used implementations

are Lloyd's/Forgy's (Forgy, 1965; Lloyd, 1982), Hartigan-Wong (Hartigan & Wong, 1979), and MacQueen's (MacQueen, 1967) algorithm.

Lloyd-Forgy Algorithm

Lloyd's, or Forgy's, clustering approach is the most well-known of the different k-means algorithms and is the simplest in its implementation. It works by first initialising a user-defined set of random data points which will act as the cluster centroids, the central positions of a cluster of datapoints. Next, all datapoints are assigned to their closest centroid and all centroid positions are updated to become the mean of all the datapoints belonging to it. This process is then repeatedly performed until a certain number of iterations or convergence is reached and results in the centroids slowly moving into the centres of different groups, or clouds, of data. A basic pseudocode of the algorithm can be seen in Algorithm 1.1.

Algorithm 1.1 - Lloyd-Forgy's k-Means Clustering

```
1:  Randomly select k datapoints to act as cluster centroids
2:  While convergence or user-defined iterations is not reached:
3:      iterations = iterations + 1
4:      for each centroid in k centroids:
5:          assign closest datapoints to centroid
6:          calculate mean value of all datapoints
7:          update centroid to become calculated mean value
8:      Determine if convergence is reached
9:      if convergence reached:
10:         end
```

While Lloyd's/Forgy's algorithm can perform clustering in an efficient manner, a common problem with the approach is that the final cluster centroids are heavily influenced by the initialisation, and often results in clusters which are not well optimised to represent global structures within the data, instead focussing on smaller regions in feature space. To address these issues, other variations of the k-means algorithm can be used.

MacQueen's Algorithm

MacQueen's algorithm shows improvements over Lloyd/Forgy's approach by better modelling the global structure of data in feature space by updating the cluster centroids every time a new datapoint is assigned, instead of recalculating only at the end of all assignments. The algorithm is largely identical to the Lloyd/Forgy approach except for this change, which is implemented by iterating over every datapoint instead of over every centroid (see Algorithm 1.2). Because the mean of each centroid is updated every time a new datapoint is added to the feature space, MacQueen's algorithm is more computationally and time expensive than Lloyd's approach. It does however offer increased clustering accuracy by relying less on the initial cluster centroid positions, although it still is common for the MacQueen's algorithm to underrepresent the

distribution of global structures in the data, thereby suffering from the same problems as the Lloyd/Forgy approach.

Algorithm 1.2 – MacQueen's k-Means Clustering

Randomly select k datapoints to act as cluster centroids
While convergence or user-defined iterations is not reached:
    iterations = iterations + 1
    for each datapoint:
        assign datapoint to closest centroid
        calculate new mean value of centroid
        update closest centroid with new mean value
    Determine if convergence is reached
    if convergence reached:
        end while loop

Hartigan-Wong Algorithm

To lessen the problems with cluster optimisation that exist due to the initialisation of the cluster centroids, the Hartigan-Wong algorithm was developed which allows for better representations of the global structure of data. This algorithm differs from the other two previously mentioned techniques, insomuch that it initialises the cluster centroids by randomly assigning all datapoints in the whole dataset to each cluster. Each datapoint is then reintroduced to each centroid and is assigned to the centroid with the smallest sum squared Euclidean distance. Once all the datapoints have been assigned, the centroids are recalculated and the process repeats until, like with MacQueen's and Lloyd/Forgy's algorithms, either a set number of iterations or convergence is reached, a representation of the algorithm can be seen in Algorithm 1.3. The Hartigan-Wong algorithm is the most computationally expensive method of performing k-means of the three mentioned in this thesis but can assign datapoints to cluster centroids more accurately than the others as the entire dataset is considered for every centroid in each iteration.

Algorithm 1.3 – Hartigan-Wong k-Means Clustering

Create k centroids
Randomly assign all datapoints to the k centroids
Calculate mean values for each centroid
While convergence or user defined iterations is not reached:
    iterations = iterations + 1
    for each datapoint:
        for each centroid:
            calculate sum squared difference between centroid and point
        assign datapoint to centroid with smallest sum squared difference
        recalculate centroid value to be mean of all assigned points
    Determine if convergence is reached
    if convergence reached:
        end while loop

1.3.2 Self-Organising Maps

Self-organising maps have already been shown to be able to successfully cluster a large set of spectral imaging reflectance data in a previous study performed by Kogou et al. (2020). One of the main outcomes of this thesis is to develop a clustering approach that can improve upon the previous implementation of the algorithm and find a general use for clustering techniques to be used on spectral imaging data of different types. Therefore, it is important that SOM is covered in detail here so that new implementations of the algorithm can be evaluated later for their application on other datasets (see chapter 2 for a detailed analysis of SOM).

The Self-Organising Map (Kohonen, 1990, 2001) is an artificial neural network-based clustering method based on the self-organisational characteristics of the cerebral cortex within the human brain. It is capable of learning directly from input data and functions through a 'competitive learning' process, where the neurons in the network directly compete with one another to represent specific parts of the data. While there are many types of neural networks that can implement competitive learning, SOM is unique in that it can represent groups of multidimensional data as clusters in topologically structured two-dimensional maps, whose structure is created through the relationships between codebook vectors, which are the vectors associated with each neuron. In the context of spectral imaging, the codebook vectors belonging to each neuron would be the mean cluster spectra belonging to each cluster and are calculated by taking a mean over all the individual constituent pixel spectra. Oftentimes in applications of SOM, the two-dimensional topographical neural maps can be used to aid in the visualisation and analysis of data, however, in the studies carried out throughout this thesis these maps are largely considered redundant, with more informative analysis instead being carried out directly on pigment maps of real spectral images, constituent cluster spectra, and mean cluster spectra grouped and produced by SOM respectively. Regardless, in all implementations of SOM, a neural network is initialised in either a rectangular or hexagonal grid or map, where each neuron will act as the centroid of a cluster to eventually create an organised map. Within the neural network map, each neuron has a weight associated with it which changes as data is introduced, with the neuron and weight being initialised by either: sampling individual data points from the input data, which can often lead to premature convergence similarly to the earlier mentioned k-means approaches such as the Lloyd/Forgy or MacQueen algorithms; or assigning random values to each weight, which can increased training times (as convergence can take longer to reach). After the weights and neural network have been initialised and built, unsupervised SOM is typically implemented in one of two ways, either in an 'online' mode or a 'batch' mode, which both correspond to different ways in which SOM assembles itself into a map and learns from the input data to create different clusters.

Online SOM

In unsupervised clustering applications, 'online' typically refers to the process whereby a model learns the best predictors for performing clustering by sequentially using a single datum at a time. In the context of SOM, this online approach involves a stochastic process of updating weights in the map every time a new datapoint is introduced into the network and assigned to its closest neuron, which is also known as the best matching unit (BMU). Typically, the BMU is found by calculating the smallest Euclidean distance between the input datapoint and each neuron, with the neuron possessing the smallest distance becoming the BMU. One of the unique characteristics of SOM is that when a neuron is updated, others within its local neighbourhood will also be changed depending on a neighbourhood function, allowing for the self-assembly, or organisation of the hexagonal or rectangular neural map. Where the neighbourhood function affecting the neurons around each BMU cab be described by:

$$h_{mn,kj}\,(t) \;=\; h(||c_{mn} \;-\; c_{kj}\;||_2^2, t) = \eta(t) e^{-\frac{||c_{mn} - c_{kj}\,||_2^2}{2\sigma^2(t)}}$$

In this neighbourhood function, c_{mn} and c_{kj} are the coordinates of the winning neuron at position m, n and another neuron at position k, j in the map respectively. $\eta(t)$ and $\sigma(t)$are the learning rate and width of the kernel respectively, and both decrease over the number of iterations, t. The neighbourhood function follows a gaussian shape and gradually reduces as the position k, j increases, until $h_{mn,kj} \to 0$ at a sufficient distance away. It is with this neighbourhood function that the topological ordered nature of the map can be created over many iterations, as it is used to update the weights of each neuron where every update can be described by the following equation:

$$w_{kj}(t+1) = \; w_{kj}(t) + h_{mn,kj}(t)\big[z_p - w_{kj}(t)\big]$$

Where w_{kj} is the weight of a neuron at position k, j, and z_p is the datapoint being shown to the neural network at iteration t. Iterations are continuously increased until a good, stable organised map is found, normally determined by a small quantisation error, which is the sum of all Euclidean distances between the cluster centroid neurons and their constituent data points. Typically, when this quantisation error remains small and does not change significantly as iterations increase, it is considered that the algorithm has reached convergence. The online SOM process can be seen in Algorithm 1.4. As it updates weights and assigns data to clusters sequentially, online SOM can be less time efficient than its batch counterpart, making it less desirable for large-scale studies. However, due to its stochastic nature, the online implementation is often proposed to be the more

stable of the two SOM modes (Cottrell et al., 2018), and therefore could outperform the batch method in terms of accuracy if the initialisation of the neural network has too much of an effect on the final clustering outcomes.

Algorithm 1.4 – Unsupervised SOM (Online)

build neural network
while convergence or user-defined iterations is not reached:
    iterations = iterations + 1
    for each datapoint:
        calculate the Euclidean distance between the datapoint and all neurons
        assign datapoint to BMU neuron and update codebook vector
        adjust weights of nearby neurons using neighbourhood function
    determine if convergence is reached or if max iterations reached
    if convergence or max iterations reached:
        end while loop

Batch SOM

As the online version of the SOM algorithm can be computationally expensive and result in slightly different outcomes due to its stochastic nature, a more deterministic batch version of the SOM algorithm was developed by Kohonen to allow for reproducible results when parameters are fixed. In this batch algorithm, clustering is not performed by updating weights after each data point is matched up to its BMU and is instead completed by updating all weights simultaneously after passing all datapoints into the neural network. The basic algorithm for this can be seen in Algorithm 1.5.

Algorithm 1.5 – Unsupervised SOM (Batch)

build neural network
while convergence or user-defined iterations is not reached:
    iterations = iterations + 1
    for each datapoint:
        Assign datapoint to nearest BMU
    update codebook vectors for all neurons
    adjust weights for all neurons and update neighbours
    determine if convergence is reached or if max iterations reached
    if convergence or max iterations reached:
        end while loop

The Batch algorithm offers many computational advantages to online SOM and has been proven to work successfully in many scenarios, however it is more sensitive to initial parameters much like k-means and can suffer many of the same problems where clusters may converge towards local structure instead of fitting for global features. In fact, when the neighbourhood function is set so that the BMU will not interact with any other neurons, effectively having a neighbourhood=0, batch SOM reduces to the same process as the Lloyd-Forgy k-means clustering algorithm (Cottrell et al., 2018).

Online & Batch SOM for spectral imaging data

While the previous algorithm used for spectral imaging showed success in using online SOM, there are many arguments that support Batch SOM to be the better performer for clustering large-scale datasets. Not only is it deterministic and repeatable, but it is also more time efficient due to having no requirement to update weights after each datapoint is introduced to the network. On the other hand, the online implementation is potentially more stable and less likely to result in suboptimal clustering results caused by neuron initialisation that will be seen more commonly in the batch mode. As it cannot be certain which of the two modes would perform better for the spectral reflectance data, it seems appropriate for each technique to be evaluated in greater detail. These evaluations are performed in detail in chapter 3, where the two different modes are tested alongside different distance calculations using the spectral imaging data recorded for the maritime Southeast Asian manuscript collection at the British Library.

1.3.3 Hierarchical Clustering

Another common deterministic clustering technique implemented in many disciplines is hierarchical clustering. The typical approach is to perform a process known as agglomerative hierarchical clustering which merges singletons (single data points) and cluster datapoints by monitoring a distance relationship between them. In this approach, datapoints are individually treated as their own cluster at the beginning of the process and then are iteratively merged, recalculating new distances after each merging, and repeating until all clusters become one. The sequence with which datapoints are grouped can be illustrated as a hierarchy of merges and is often represented as a dendrogram, acting as a hierarchy tree, which can then be cut in different ways to gather a final group of clusters. Alternatively, a divisive clustering method can also be performed which follows the opposite path to the agglomerative approach by beginning with all datapoints placed into a single group before iteratively separating the most different datapoints until every individual datapoint is its own cluster. A basic representation of the agglomerative clustering process can be seen in Algorithm 1.6.

Algorithm 1.6 – Agglomerative Hierarchical Clustering

create K clusters from N datapoints where K=N
while total clusters, K > 1:
    calculate distances between all N datapoints
    find closest matching pair of datapoints
    merge closest pair into a new datapoint with new value
    remove merged datapoints and add new datapoint to pool so N = N - 1
    If total clusters K = 1:
        end while loop

Linkage Criteria

When calculating distances during any form of hierarchical clustering, individual singleton datapoints can be compared by simply computing the distance between the two points. However, it is important that when singleton or cluster datapoints are grouped with other clusters, the data used during comparison can be controlled for different purposes. In hierarchical clustering, these positions can be defined by selecting different linkage criteria which change the part of the cluster which is compared. There are many different linkage criteria that can be selected, but the most common include:

- Single Linkage, where the closest constituent data points within different clusters are compared.
- Complete Linkage, where the furthest constituent data points within different clusters are compared.
- Average linkage, where the average distance is computed from the distances between all data points within different clusters before being compared.
- Centroid linkage, where cluster centroid positions are computed and compared for different clusters.
- Ward's Linkage, where the two clusters to be merged will be those with the minimum between cluster sum of squares distance.

No single approach is more advantageous over all applications, therefore when choosing an appropriate linkage method, the main goal of the clustering and its theoretical justifications should be considered, though for most applications which use hierarchical clustering purely to create groups of similar 1D data, Ward's linkage tends to be used most, such as in (Gumansalangi et al., 2023; Hashemi-Nasab & Parastar, 2022; R. Kumar & Sharma, 2017; Moros et al., 2008)

Dendrogram Trees

The output dendrograms showing the hierarchical structure of the clustering process can be visualised and cut to acquire different numbers of clusters from the data. Typically, clustering dendrograms will be represented as column graphs, illustrating the order with which different datapoints and clusters merge and representing the distances between clusters as heights. Therefore, the more similar the two clusters are, the smaller their heights will be in the dendrogram. The dendrogram can be cut to acquire clusters at different stages of the hierarchy, where cutting at a specific global height, h, will produce k clusters, one for each intersection between a "tree branch" and the height. An example of a basic dendrogram can be seen in Fig. 1.6, with lines illustrating the cutting positions at different heights and the number of clusters produced as a result.

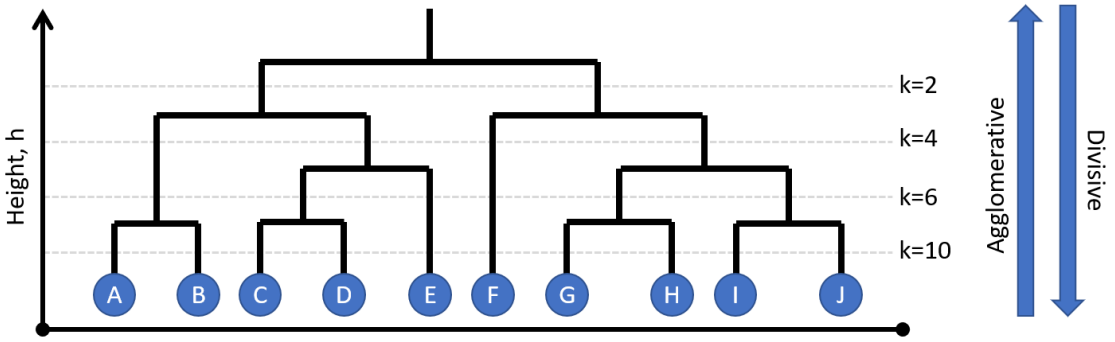

Fig.1.6 Example of a dendrogram tree commonly encountered in hierarchical clustering applications. Notice that cutting the tree at different heights will produce varying numbers of clusters, k.

Hierarchical Clustering of Spectral Imaging Data

Hierarchical Clustering can offer many useful advantages over other clustering techniques when applied to clustering spectral imaging data, such as its ability to easily implement different linkage types when computing the similarity between datapoints or clusters. It also offers a convenient way of visualising how clusters with high dimensionality relate to each other through the creation of dendrograms. However, hierarchical clustering is computationally expensive when dealing with large amounts of data, as the storage requirements for performing hierarchical clustering increase exponentially as the number of datapoints increases; a result of having to create large dissimilarity matrices to identify the relationships between different datapoints.

In large collections, performing hierarchical clustering on entire sets of spectral imaging data would require huge amounts of computer memory and storage, making the application of hierarchical clustering impractical for real studies. For example, when performing hierarchical clustering of spectral imaging data cube containing 1 million spectra, a dissimilarity matrix with a total size of 7.28 terabytes would be required for double precision data, which most computers would not be able to access or use efficiently, even if the data is stored externally and not held in RAM. Therefore, if hierarchical clustering is to be used, it would be more advantageous to use it for more precise and deterministic clustering after other techniques, such as k-means or SOM, have been used to reduce the size of the dataset in need of clustering.

1.3.4 Clustering Evaluation

To understand whether one method of clustering is better to use than another for a set of new data, it is important that some form of cluster validation can be performed to properly monitor the accuracy and stability of the technique. Cluster validation is performed either internally, where internal evaluation uses the input data alone and is useful for unsupervised clustering approaches,

or externally, where external evaluation uses outside information such as the ground truth or labelled data to determine the validity of the results. As the clustering techniques to be used throughout this thesis are all unsupervised and no labelled data exists until after the analysis of a very large collection is complete, internal validation techniques should be considered the better option to use to monitor the accuracy of different approaches on new spectral imaging datasets. For internal validation, clustering techniques are always evaluated by monitoring two main components, their cohesion and separation. Cohesion evaluates the distances between clustered input data points and their associated cluster centroids. Conversely, separation evaluates the distances between the clusters themselves. Typically, a clustering result is considered good if its cohesion and separation are high, such that the intra-cluster distance (between data points and cluster centroids) is low, and the inter-cluster distance (between other clusters) is high. Cohesion and separation of clusters can be measured by using the sum of squared errors (SSE) and the between group sum of squares (SSB), which both lend themselves to algorithms such as k-means and SOM well as these commonly use SSE during the clustering procedure to optimise the clusters themselves. Furthermore, the SSE and SSB can be collated to gather the global cohesion and separation amongst all clusters and across all variables for an entire data set, alternatively known as the total sum of squared errors (TSSE) and total between group sum of squares (TSSB) respectively.

For monitoring the accuracy of clustering approaches using spectral reflectance data, the TSSE can be calculated by using the following equation:

$$TSSE = \sum_{\lambda_{min}}^{\lambda_{max}} \sum_{k=1}^{k_{max}} \sum_{i \in k} (R_{i,\lambda} - \bar{R}_{k,\lambda})^2$$

where $\bar{R}_{k,\lambda}$ and $R_{i,\lambda}$ are respective the mean reflectance spectrum and i^{th} reflectance spectrum for different wavelengths, λ, and for each cluster k. And similarly, the TSSB can be calculated using:

$$TSSB = \sum_{\lambda_{min}}^{\lambda_{max}} \sum_{k_i=1}^{k_{max}} \sum_{k_j=1}^{k_{max}} (\bar{R}_{k_i,\lambda} - \bar{R}_{k_j,\lambda})^2$$

where $\bar{R}_{k_i,\lambda}$ and $\bar{R}_{k_j,\lambda}$are the mean cluster spectra for the ${k_i}^{th}$and ${k_j}^{th}$ clusters. Using the TSSE and TSSB as separate metrics for evaluating each clustering technique allows for a basic understanding of which approach is best; if a value has a low TSSE and high TSSB, then it can be considered better than that which has a high TSSE and low TSSB. Due to its simplicity and simple interpretation, a comparative study using TSSE and TSSB can be convenient enough to test between different clustering algorithms, however, there are many techniques which can go a step further to unify the two metrics and allow for the evaluation of clusters using a single value, though these are not implemented in this thesis.

1.4 Supervised Classification Techniques in Machine Learning

In contrast to the previously mentioned clustering techniques which can be implemented to accelerate the process of performing large-scale spectral imaging data analysis, other supervised machine learning techniques can often be implemented to perform identification of data which has already been encountered during previous scientific studies. These techniques are normally referred to as classification techniques, and unlike unsupervised methods do not learn directly from the data provided to them. Instead, these supervised classification techniques function by learning ways to represent labelled data as a set of classes, where the goal is to create classes, or groups of data, which are known to correctly represent the ground truth for different datapoints or observations. Typically, once a representation of these labelled datasets has been learned by a model, normally referred to as training, the trained classification technique can be used to automatically assign new unseen datapoints into the created classes. This is basically performed by comparing the trained classes with these new observations and calculating which class is the most similar and therefore the class to which the new observation should belong.

In the context of this thesis, this implies that one may use a classification method to learn a representation of a dataset with known artistic materials so that new unseen and uncharacterised datasets can be automatically classified. This could apply to many different materials so that representations of specific paper, inks, pigments, or pigment mixtures, could be learned and then used to classify illuminated manuscripts. If the studies performed throughout this thesis can provide brand new ground truth data for materials via the large-scale holistic study of pigments in maritime Southeast Asian manuscripts, then newly collected data in the future may be able to be automatically classified into previously identified pigment classes. Many different classification techniques have been developed in the past and used to attain certain research goals in different studies, however, the number of classification algorithms available today is vast and so their implementation for the goals in this research is not clear. In chapter 5 of this thesis, a handful of these different classification models are evaluated for their use in automatically identifying different materials and pigment mixtures for studies of illuminated manuscripts. The techniques implemented are therefore covered in greater detail here for a better understanding of the different models.

1.4.1 Supervised Self-Organising Maps (SOMs)

While self-organising maps were covered in detail in the previous section for their use in unsupervised clustering, SOMs can also be used in supervised ways (Kohonen, 1990, 2001), a method that was previously used as part of the original ISAAC lab's clustering code developed by Kogou et al., (2020). This supervised application however also means that in addition to clustering, SOM can also be used for classification purposes, where a self-organising map can be

used as a trained model by using a set of labelled data. For SOM, this specifically involves using the unsupervised method to train the neural network to learn a representation of the input (training) data which creates an organised map of neurons which can be used for classification. In a perfectly trained network, each individual neuron would correspond to a specific class of labelled data, with the individual neurons being arranged in the map with more similar classes positioned closely together, however, there are many applications where networks larger than the number of classes are used to uncover underlying structure in the labelled data and improve the accuracy of fitting new data, such as in Wong et al., (2019) where a massive SOM is used to improve classification accuracy. During classification, each neuron in the SOM will be associated with a single value or vector, normally the codebook vector, which is used to compare against new input data to find similarities and perform classification, in much the same way that the BMU is used during the unsupervised approach. As supervised SOM has already demonstrated its use in VNIR studies with data like that seen in this thesis, it makes sense that performing classification using it would be recommended in this research project.

1.4.2 **Spectral Angle Mapper (SAM)**

Another commonly implemented classification technique is the spectral angle mapper (SAM) (Kruse et al., 1993). The SAM typically performs classification by comparing reference spectra, also commonly referred to as "endmembers", with new unseen data. This is performed in SAM by calculating a spectral angle-based similarity in feature space between the reference spectrum and the target spectrum being classified. This angle is typically calculated for every class within a model, where the reference with the smallest angular distance from the target will dictate the class to which the new spectrum belongs. Like with supervised SOM and other classification techniques, the SAM requires training data to be able to function and has shown to perform well in classifying spectral imaging data of different kinds.

In the cultural heritage community, SAM is regularly used to perform pigment identification by using known pigment endmember spectra to map the distribution of pigments in a spectral image (Grabowski et al., 2018). This makes the technique potentially very useful for further analysis performed after identifying materials in the research within this thesis, however, the stability and accuracy of the technique is not known for all spectral imaging instruments and types. Furthermore, there are very few applications of the SAM which highlight its use over thousands of spectral images in a single study. As it is therefore not completely known how well it may perform in the context of this thesis, the technique is tested in chapter 5 along with other classification models.

1.4.3 K-Nearest Neighbours (k-NN) & Nearest Centroids

Another commonly implemented classification technique, k-Nearest Neighbours (Cover & Hart, 1967; Fix & Hodges, 1951) is one of the simplest classification models and is often used in many different applications, though less so due to the rise in use of other classification techniques. It primarily works off the understanding that most similar datapoints or vectors, in this case, spectra, can normally be found nearby to one another within feature space, with the most alike being considered neighbours. k-NN effectively performs classification by placing a target datapoint into feature space which has already been trained on an initial set of data. Because the datapoints within the training dataset have known classes, the target can be associated with a set of nearest neighbour datapoints in feature space that have their own class values. A free parameter, k, can be used to define the number of neighbours typically used to identify a class, however, regardless of the number of neighbours a class is always picked via a majority vote. This implies that the most common class found amongst the nearest neighbours of the target datapoint will be the class to which the target is assigned.

Nearest centroids is a similar algorithm to k-NN, however instead of dealing with individual data points in feature space and performing a majority vote to perform class assignment, distances are instead computed between target datapoints and a cluster centroid. The cluster centroid is normally computed by calculating the mean of all the datapoints belonging to the class in feature space and is similar in nature to a typical k-means clustering algorithm insomuch that target data is assigned to the cluster centroid (and therefore class) which is the smallest distance away.

1.4.4 Multilayer Perceptron

The multilayer perceptron (MLP) is a feed-forward neural network-based method which normally consists of an input layer, output layer, and one or more hidden layers, where each layer is a set of neurons connected to the next. MLPs perform classification by sending training data through the input layers with known outputs. During training, the hidden layers contain weights that adjust through backpropagation (Rojas, 1996) to learn a representation of the data which can classify new unseen target datapoints. Their use in classification is well documented in different fields, but in the cultural heritage community MLPs have seen limited use introducing the question of whether they will be useful for the research carried out in this thesis. Some examples of their use in other applications, however, can be seen in Grabowski et al., (2018), Lin et al., (2019), Babić, (2022).

1.4.5 Support Vector Machines (SVM)

The support vector machine (Cortes & Vapnik, 1995), or SVM, is another supervised learning technique which can be used to perform classification on different types of multidimensional data and has been shown to perform well in classifying spectral imaging data for artworks and pigments in cultural heritage studies. Furthermore, in many comparative studies, SVM is regularly shown to outperform other techniques such as supervised SOM (Chaplot et al., 2006; Salouan et al., 2014) and k-Nearest Neighbours (Thanh Noi & Kappas, 2017). The basic idea behind the support vector machine is to use labelled data to train vectors which can be used 'cut' feature space into different sectors where each sector roughly represents a single class. More specifically, these vectors, known as hyperplanes, essentially act as the decision boundaries between different class areas and can be produced because of any function (or kernel), with some common implementations involving linear, polynomial, and radial-basis function kernels, all of which can result in different classification accuracy and outcomes. Because of the use of these different kernels, SVM has shown a propensity to solve nonlinear problems, often making it more beneficial than other techniques. In studies using spectral imaging data, SVMs have shown to be effective in performing classification, some examples of which can be seen in (Polak et al., 2017), (Gao et al., 2023).

1.4.6 Decision Trees & Random Forest

Decision trees are a branch of non-parametric machine learning methods which can also be used for performing classification tasks and have demonstrated their use in classifying pigments in the past (Fan et al., 2018). Like the previously covered techniques, they also function using supervised learning, but instead perform classification by attempting to segregate data into basic decision rules. This implies that a decision tree classifier can essentially be broken down into a large collection of binary choices which spreads from an initial origin, creating a "tree" of decisions after training which can be used to assign new data into classes. A common problem encountered with decision tree classification however is that they are vulnerable to overfitting.

In response to the overfitting problems, the random forest was developed as a unique approach to using decision trees, where instead of using only a single tree to assign data into classes, many different random decision trees, assembled into a "forest" can be used. Random forests have been shown to provide increased accuracy in many classification applications, and therefore show promise when performing classification of spectral imaging data.

1.5 Thesis Structure

The interdisciplinary nature of the research carried out in this thesis means that there is a spread of different topics explored throughout the different chapters, though the overarching theme remains with how machine learning and spectral imaging can be used for studies in cultural heritage.

Chapter 2 begins by first performing an evaluation of unsupervised clustering techniques on reflectance spectral imaging data, where the previously discussed k-means and SOM are investigated in detail. The chapter then continues to discuss some of the difficulties with performing clustering in the VNIR and how implementing both spectral and colour information directly into a novel clustering method can improve accuracy when clustering any dataset. This new methodology is subsequently used to automatically cluster the British Library's maritime Southeast Asian illuminated manuscript collection, after which a new software-driven grouping procedure is then introduced and demonstrated on a sample of the collection, where the highly accurate clustering results can be further grouped based on different research questions. Chapter 3 follows almost directly from chapter 2 and delves into the holistic multimodal analysis of the British Library's maritime Southeast Asian illuminated manuscript collection. The chapter begins with a brief overview of the collection and is proceeded by the analysis of artistic materials, namely pigments and inks, used in many different manuscripts. Throughout the chapter, grouped pigment and pigment mixture maps are created alongside spectroscopic analysis to illustrate the distribution of pigments and mixtures so that confidence can be gained in pigment maps being accurate representations of the ground truth. As pigment mixtures are identified and artistic practises uncovered, discussions into the likely trade connections between maritime Southeast Asia and the rest of the world are considered. The clustering and grouping approach is then discussed once more with some further pitfalls being demonstrated before a final collection-wide discussion is made to summarise the findings of the chapter.

As mentioned, chapter 4 introduces a novel methodology for performing multimodal clustering in numerous spectral domains. It begins first with demonstrating how the clustering method developed in chapter 2 can be used for clustering XRF mapping data after implementing a pre-processing procedure. Following this, a common theme when performing clustering becomes apparent across both XRF and VNIR: that clustering in only one spectral domain at a time fundamentally limits the accuracy of performing holistic studies of artistic materials. In response, the clustering method introduced in chapter 2 is adapted so that multimodal clustering can be extended to allow for any number of techniques to be used. The developed method is subsequently used on Peruvian and Chinese export paintings to investigate the use of their materials using VNIR, XRF and ER-FTIR mapping, where it is found that this new holistic multimodal clustering not only improves accuracy but can also reveal new information about paint layering.

Chapter 5 evaluates the uses of classification techniques for performing automated pigment mapping using results found in chapter 3. In this chapter, many different commonly implemented classification techniques are tested, after which the best performing approach is implemented on the British Library's maritime Southeast Asian collection to investigate the validity of using clustering to train classification models.

Then finally, chapter 6 briefly discusses a new spectral imaging and machine learning-based analysis software application, known as Guisi, used throughout this thesis and in multiple other studies for the visualisation and interrogation of complex spectral imaging data and machine learning clustering results.

Chapter 2
Automated Clustering & Grouping of Large Spectral Imaging Surveys

2.1 Introduction

To collect large amounts of data useful for identifying artistic materials over entire collections, spectral imaging in the visible-near infrared (VNIR) is often implemented as it can efficiently provide vast amounts of detailed chemical information about the use of pigments and materials present at different regions on an object. When performing a study which is both large in volume and variety, it is important to collect data for a large enough number of items in a collection so that statistically reliable conclusions can be made about the types of materials used. However, as the amount of data being accumulated grows it becomes more difficult to analyse in traditional ways, which has necessitated the active development and use of automated techniques to aid in the processing and analysis of large amounts of spectral imaging data. Today, most of the techniques often implemented to aid in the automation of spectral imaging data analysis are centred around the implementation of statistical or machine learning algorithms, such as PCA (Rodarmel & Shan, 2002), t-SNE (Alfeld et al., 2018; Melit Devassy et al., 2020; Melit Devassy & George, 2020), k-means clustering, and others (Fiorucci et al., 2020; L. Liu et al., 2023). However, at the time of writing this thesis, there are very few well-functioning approaches which utilise these techniques in a general and automated way for tackling the difficulties involved with analysing whole large collections. Instead, most implementations are normally focussed on specific problems, require significant human parameterisation, transformation into different feature spaces or supervision, or need a priori knowledge about the data to be performed efficiently. Therefore, there is a desire for the development of a general clustering algorithm designed to be implemented for large spectral imaging surveys for collections in cultural heritage.

There only clear prior study for performing clustering in this way can be found in Kogou et al., (2020) which was able to achieve unsupervised clustering for many spectral data cubes on wall paintings, which functioned well for their project at the time. However, it was found that when expanded to new datasets, problems were encountered with this approach where misclustering of data became common for materials with similar spectra or high absorption, even if the colour was noticeably dissimilar which typically implies a material difference. This thesis, therefore, seeks to redevelop and improve upon this previous approach, to rectify the problems and create new, more precise clusters in an automated way. This chapter evaluates machine learning approaches in their use on spectral imaging data collected in the VNIR and introduces a novel method of performing clustering for large collections based on the combined use of both colour and

reflectance. Furthermore, a novel grouping methodology is then also presented which allows for the merging of highly precise clusters into different levels of complexity based on what the desired research questions of any large-scale study may be.

2.2 Instruments & Data Collection

The bulk of the data collection for the study of the maritime southeast Asian collection was performed using VNIR reflectance spectral imaging so that large amounts of data could be captured for performing material and pigment identification. Throughout the study, two different spectral imaging systems were used, PRISMS (Liang et al., 2014) and MegaVision EV (Duffy, 2018; MegaVision, 2023), however, most data collection was performed using the former, and this was the only system which was fully evaluated using the clustering techniques outlined in this chapter.

PRISMS Reflectance Spectral Imaging

The PRISMS spectral imaging system has a well-documented history of use in cultural heritage studies and has demonstrated its usefulness for large-scale data collection where it has been used for the study of wall paintings (Kogou et al., 2020). More importantly for this thesis however, it has also demonstrated usefulness in performing pigment identification studies of paintings on paper (Kogou et al., 2015, 2016), and due to its ability to perform this data collection in a largely automated way, it stands to reason that PRISMS would be an excellent candidate for the large-scale study of an illuminated manuscript collection.

In the setup for collecting data in the VNIR, PRISMS utilises a basic filter wheel system, where reflected light from an object is passed through a lens and then directed through a wheel with 10 different optical filters, before being captured by a CCD detector. For all the data collected with PRISMS in this thesis, the 10 optical filters used ranged from 400nm to 850nm, with regular steps of 50nm and with each filter possessing a bandwidth of 40nm. The output of PRISMs after processing is therefore a 10-channel deep spectral reflectance data cube where each pixel in the 2D image plane possesses a spectrum ranging from 400-850nm, as demonstrated in Fig. 2.1. For each spectral data cube, a colour image is normally also produced and is created directly from the reflectance information, allowing for pixel-perfect colour-spectral comparisons to be performed during analysis.

Most of the data capture for PRISMS is automated and the system can be run in a scanning mode so that entire paintings or manuscripts can be recorded from a fixed position. The images produced by the system are 1360x1024 pixels in size and can achieve a spatial resolution of less than 100μm at distances of less than approximately 3-3.5 metres. The multiple images can be

mosaiced in post-processing to combine scanned images into a large single mosaic (see Fig. 2.2), also making PRISMS a useful tool for large-scale data visualisation.

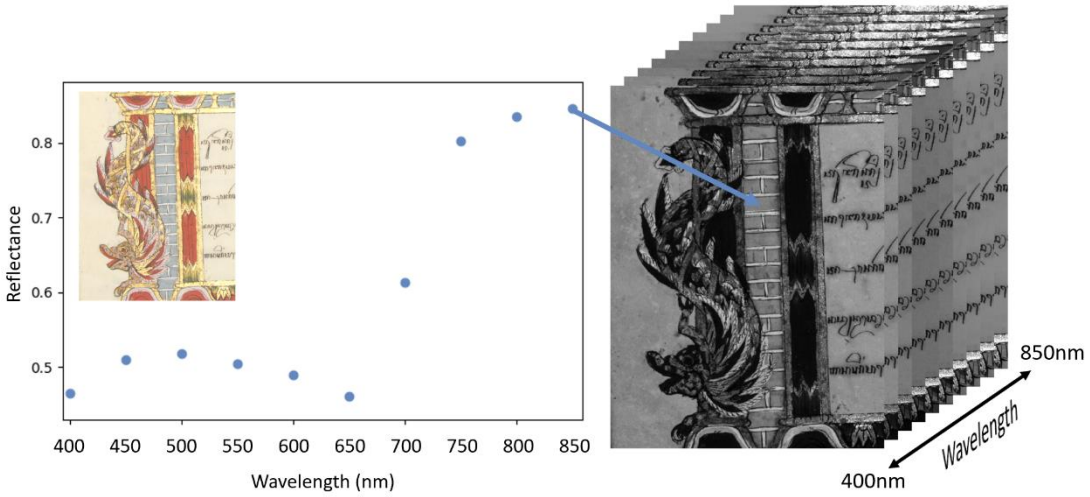

Fig.2.1 A basic representation of a PRISMS spectral imaging data cube with wavelength 400-850nm. An extracted spectrum for a single pixel can be seen on the left-hand side.

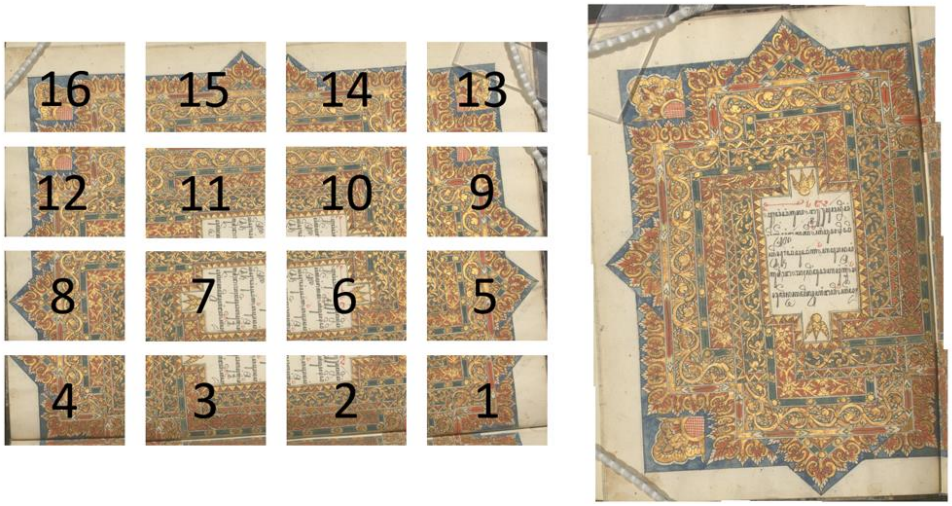

Fig.2.2 Example of a mosaiced PRISMS scan, where 16 different images from manuscript MSS Jav 24, folio f2v, of the British Library's collection are combined to create a single larger image.

PRISMS can be used to record large amounts of data quickly, as it can take only a few seconds to record a single spectral imaging data cube and with most of the process being fully automated, large amounts of data can therefore be collected quickly, however, this comes at the cost of lower spectral resolution compared with other methods such as diffraction grating based techniques. This lower spectral resolution could be considered a cause for concern when identifying certain materials, however, as most pigments in the VNIR range normally only possess broad spectral

features, a higher spectral resolution is not always needed (Liang, 2012). Furthermore, PRISMS is still normally able to distinguish some pigments with these fine absorption features regardless, such as in the identification of red dyes, though in these scenarios the specific type of dye used may remain ambiguous. The system is therefore useful for providing large-scale surveys of different materials, allowing for many different pigments to be distinguished apart from one another quickly, and even when more conclusive identification of certain materials is required using alternative techniques or higher spectral resolution, there is still an advantage provided as analysis can be guided more clearly by PRISMS.

Data Collection at the British Library

Spectral imaging data for the British Library's maritime Southeast Asian manuscript collection was recorded on-site in illumination-controlled conditions with two halogen lights illuminating the manuscripts and calibration targets from approximately 1 metre away. Manuscripts were placed to ensure that the plane of imaging was as perpendicular to the optical axis of PRISMS as could be made possible, ensuring that the entire FOV of an image would remain in focus during scanning. In total, spectral reflectance data was recorded for 50 illuminated manuscripts and involved the capture of 1515 data cubes, resulting in almost two billion spectra after processing.

In typical studies, these two billion spectra may be quickly interrogated to investigate whether the material composition is roughly the same for different colours throughout an entire spectral data cube or manuscript. However, when dealing with many spectrally complex items and over a thousand spectral data cubes, the assumptions made in analysing data in this way become more error-prone, as many important features could be missed by monitoring data only by eye. This therefore clearly highlights the need for automation using machine learning techniques when analysing materials throughout the entirety of the British Library's maritime Southeast Asian manuscript collection.

2.3 Clustering Tests for PRISMS Spectral Reflectance Data

With PRISMS producing such large amounts of spectral imaging data, computational methods of dealing with such large datasets had to be investigated, to ensure that analysis could be performed in a timely and realistic way. To that end, machine learning offered the ability to reduce the volume of data down to a more manageable dataset using clustering. At the beginning of this research, one clustering approach for analysing large amounts of spectral imaging data using PRISMS had already been developed for the analysis of wall paintings, as seen in Kogou et al., 2020). The clustering approach outlined in this previous investigation had performed well for the research questions proposed, however when expanding the method to alternative datasets the algorithm suffered from multiple difficulties where misclustered spectral data, meaning spectra which were incorrectly clustered into non-representative groups, became a common issue for

highly absorbing materials and spectra with similar shapes and intensity. Furthermore, due to the implementation of supervised self-organising maps (SOMs), the procedure could become very computationally expensive and time-consuming. To address these problems a wholly unsupervised clustering technique which could be used for multiple different large and complex datasets, like the manuscript collection involved in this study, was therefore required.

2.3.1 Preliminary k-Means and SOM Testing

To develop a new clustering approach for large-scale data analysis, a preliminary investigation into the capability of different clustering techniques had to first be performed. Two techniques were chosen for testing on PRISMS data, namely k-means and unsupervised SOM. These were chosen as they offered the least computationally expensive approaches to performing clustering for large-scale datasets. In addition, they also can be run with little parameterisation other than the number of clusters, making them simple to use. However, to understand the best setups for comparing each technique, different algorithms, which had not been previously tested, were evaluated prior to comparing k-means and SOM for large collections.

Test Spectral Imaging Data

A spectral imaging cube taken from folio f104v of manuscript MSS Jav 24 from the British Library was selected to act as the test dataset for the different algorithms, which after processing measured 1292x962 pixels spatially, meaning that the total number of pixels to be clustered is 1,242,904 where each possesses a 10-band reflectance spectrum ranging from 400-850nm.

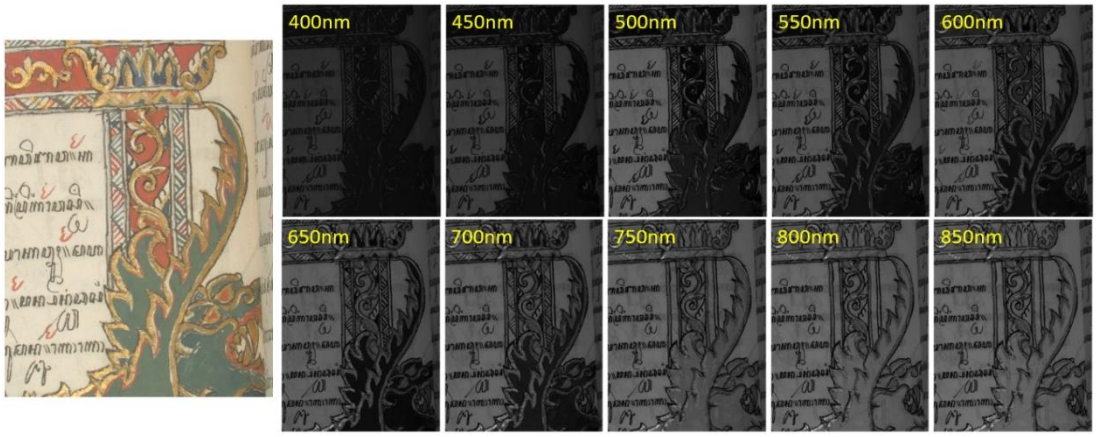

Fig.2.3 Representation of the PRISMS test data set used for comparing different clustering algorithms. Each greyscale image depicts a single wavelength band from 400-850nm.

k-Means Algorithm tests

As mentioned in chapter 1, k-means clustering has three main algorithms that are commonly implemented to perform clustering: Lloyd/Forgy's Algorithm, MacQueen's Algorithm, and the Hartigan-Wong algorithm. To conclude which would perform the best with PRISMS spectral reflectance data cubes, each algorithm was used to perform clustering on the test data. A total number of clusters ranging from 1-1000 were produced, with the mean total sum of squares error (TSSE) being calculated across different clusters to provide an insight into the accuracy of each approach. As k-means clustering can be sensitive to the initial cluster centroid positions with which each cluster begins, the test was run over 100 randomly generated seeds, so that any randomly under or overperforming runs would have a minimal effect on the mean TSSE. The 100 different seeds were randomly generated based on the date and time and were then fixed for all setups to be used to initialise the cluster centroids. For each technique, the number of maximum iterations was set at 100 so that the clustering would eventually halt if convergence was not reached. The results of the cluster validation, through monitoring TSSE, clearly show that the Hartigan-Wong (KM-HW) algorithm outperforms both the Lloyd/Forgy (KM-L) and MacQueen (KM-MQ) algorithms, as demonstrated in Fig. 2.4. This can be seen from both the table and graph which shows that while all three k-means approaches do eventually reach a position of convergence, the KM-HW approach performs with a lower mean TSSE than every other algorithm.

Clusters	Mean TSSE After 100 random seeds		
	KM - HW	KM - L	KM - MQ
1	4989.317	4989.317	4989.317
10	583.2074	583.453	583.569
25	314.758	317.239	316.352
50	214.9061	219.475	219.578
100	154.6257	158.484	158.603
200	112.5453	118.058	118.301
300	92.76073	99.237	99.183
400	80.43841	86.837	86.808
500	71.78408	78.353	78.177
750	57.10111	63.871	63.897
1000	47.55451	54.870	54.845

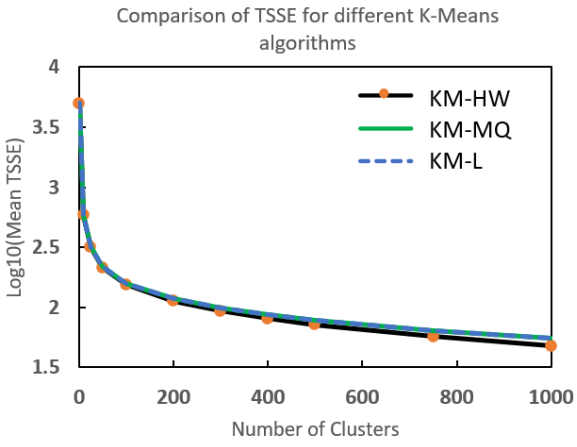

Fig.2.4 Table and a graph showing how the TSSE changes for k-means clustering algorithms as the number of clusters is increased, where K-Means with the Hartigan-Wong algorithm shows the best performance.

SOM Algorithm Tests

SOM algorithms were also investigated by evaluating the online and batch modes, where the neural map learns either iteratively over every datapoint or as a batch over the entire dataset respectively. To do this, the two different algorithms were compared using the same test used for evaluating the k-means approach, but instead, the online and batch SOM were used to perform clustering from 1-1000 clusters, under the 100 random seeds. The results of this preliminary test can be seen in Fig. 2.5. Like the results seen for the k-means test in the previous section, the SOM test shows that as the number of clusters is increased, both techniques move towards points of convergence. However, it can be noted for the SOM tests that the Batch SOM approach outperforms the Online SOM method, suggesting that if SOM is used for the clustering of large collections, the batch approach is preferable.

	Mean TSSE After 100 random seeds	
Clusters	SOM Online	SOM Batch
1	5015.93	4989.317
10	584.263	584.19
25	313.482	325.713
50	218.321	221.161
100	162.269	161.534
200	122.254	116.355
300	107.611	97.369
400	99.272	85.438
500	94.316	76.415
750	87.322	62.513
1000	82.697	52.538

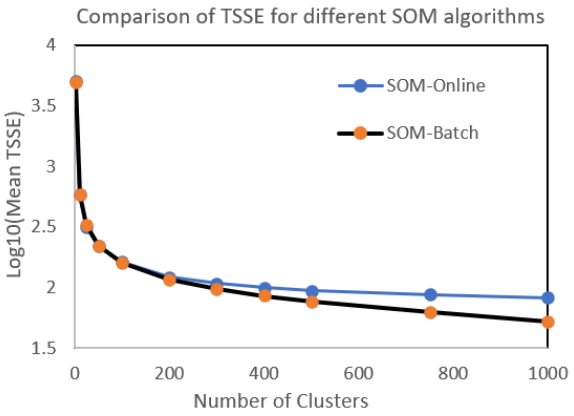

Fig.2.5 Table and a graph showing how the TSSE changes for different SOM algorithms as the number of clusters is increased, where the best performing technique was found to be the batch algorithm.

k-means (Hartigan-Wong) vs SOM (Batch)

With the best algorithm being determined for each clustering method, a further investigation was then performed comparing batch SOM and Hartigan-Wong k-means. To ensure the widest variety of data, the two methods were compared using the entire collection of 1515 PRISMS spectral imaging data cubes recorded at the British Library. The TSSE was subsequently recorded for each cube to determine which approach would have the smallest overall inter-cluster distance, and therefore the highest accuracy. The distribution of different TSSE values for the two approaches can be seen in a histogram shown in Fig. 2.6. Though the two distributions in the histogram are similar, the peak of the curve for Batch SOM has a slightly smaller TSSE than the k-means approach using the Hartigan Wong algorithm, suggesting that most Batch-SOM clustered data cubes perform better than the k-means due to them having lower TSSE values. In addition to this, the total TSSE across all spectral data cubes, calculated by adding together all TSSE values across

all cubes, is noticeably smaller for the Batch SOM and there appear to be fewer clustering outliers where the TSSE is greater than ~15000, suggesting both higher accuracy and greater stability in using SOM instead of k-means. With this result in mind, it would be appropriate to assume that the best approach to use for PRISMS spectral imaging data would be the Batch SOM algorithm so long as convergence on local features does not occur during clustering, though there is little evidence to suggest that this does occur.

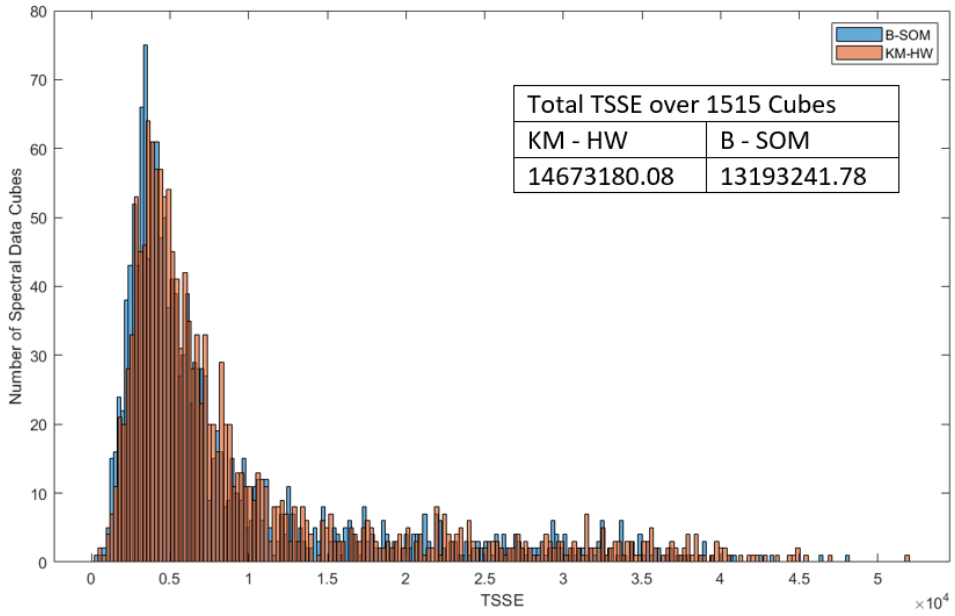

Fig.2.6 Histogram showing the performance of the Hartigan-Wong k-means (Red) and Batch SOM (Blue) algorithms on all 1515 PRISMS spectral imaging data cubes. On average, the k-means performs worse with a larger TSSE, with the total TSSE across all 1515 data cubes also being much higher than for SOM.

2.3.2 Investigation into Misclustered Spectra

When performing the previous clustering tests on PRISMS spectral imaging data, it was noticed that misclustering can occur when the number of clusters is too small, if the data possesses highly absorbing spectra (low SNR), or if differences in spectral features are small and therefore difficult to distinguish automatically. This can be easily demonstrated in the test spectral imaging dataset where two separate colours, a dark green and blue, are often clustered together due to their similarity in spectral reflectance, even though they possess a small deviation in intensity around 400-450nm (see Fig. 2.7). Both regions have a large absorption feature indicative of indigo from approximately 650-850nm and possess similar spectral shapes from 500-650nm. The only significant difference between the two spectra is a minor difference at around 400-450nm, where the intensity of the spectrum extracted from the green region is slightly lower than that of the blue, due to the additional absorption caused by the presence of a yellow pigment to create green. These small spectral changes do not appear to be significant enough differences to avoid

misclustering problems at varying numbers of clusters, however, several routes can be followed to address the issue and minimise the number of incorrectly clustered pixels. One approach would be to produce an even larger number of clusters to allow for a separation between the two regions. This could result in a huge number of clusters for other areas without fully solving the blue-green misclustering however. Furthermore, even if the number of clusters is increased there can be no guarantee that different colours with similar spectra would be fully separated over many different spectral cubes in the collection, as there is no way to easily direct the clustering to achieve the desired result using spectral reflectance alone.

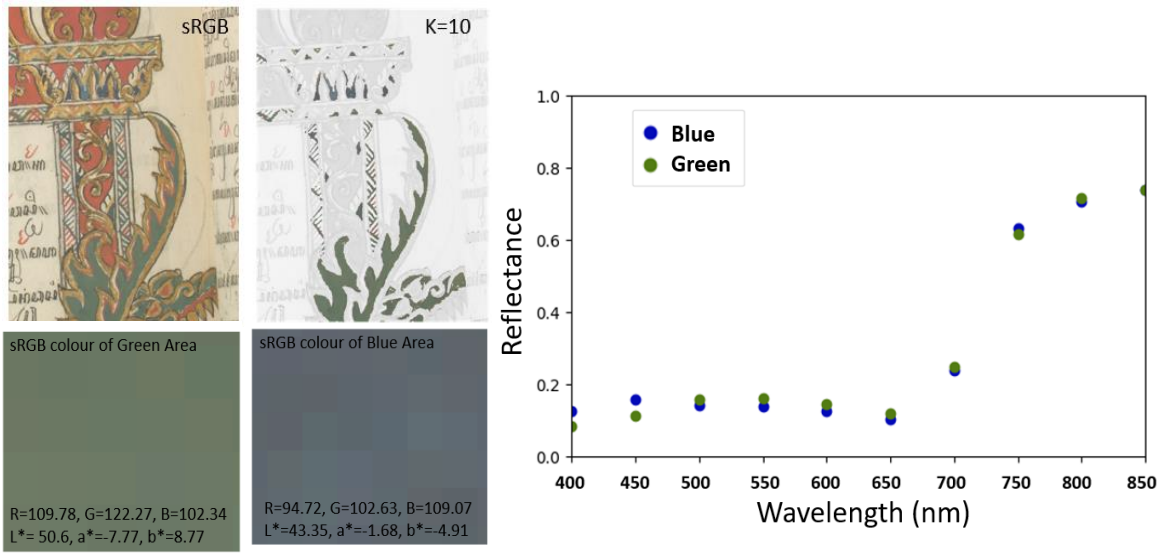

Fig.2.7 Representation of misclustered blue and green colours within the test data cube for varying numbers of clusters. A comparison of their colour and spectra can be seen on the bottom left. Shown on the right are blue and green spectra with similar features, indicating why clustering places the different colours together.

For small spectral changes such as between the blue and green mixtures, visual inspection of clusters is often performed by eye, where colours are monitored to test if spectral features are the same. Instead of performing this manually, the use of this check prompted the development of a method which can include colour information in the clustering approach itself, allowing for a separation between the blue and green areas. If using colour, many different colour spaces (e.g. sRGB, RGB16, CIELAB) could be used to resolve the issue of misclustering, however, to understand which may perform better, a comparison of SOM's performance on different colour spaces was first necessary. Two colour spaces were investigated as candidates to improve clustering results: sRGB and CIELAB:

- sRGB: is well represented in colour images which can be visually inspected in any digital format. Furthermore, the information required to form a separation between blue and green materials is split into two of three overall channels, increasing both signal-to-noise and total variance in comparison to spectral reflectance.

- CIELAB: also provides improved signal-to-noise and greater variance between indigo and indigo-yellow green mixtures in comparison to spectral reflectance. But in addition, CIELAB is the international standard colour space and is more perceptually uniform than sRGB. Furthermore, it also has a well-defined colour difference metric, ΔE^*_{ab}, which can be used to define material differences based on the human perception of colour. If we consider CIE76 as a good enough estimation of colour when used alongside reflectance, ΔE^*_{ab} can be used directly for clustering because it is calculated from the Euclidean distance between two CIELAB colour vectors. Therefore, SOM, when used on CIELAB data, is effectively performing clustering in ΔE^*_{ab} feature space.

Both colour spaces can be produced directly from the reflectance spectral data, and each can offer a different interpretation of colour, each with its own advantages, making it unclear which would be better to use in the clustering procedure. To conclude which colour space would be the best to incorporate into the process, regions of interest for the blue and green areas on the test data cube, illustrated by the white selected areas in Fig. 2.8, were selected so that the number of pixels shared between the two regions could be monitored under different clustering parameters.

SOM was performed 20 times on the sRGB, CIELAB, and PRISMS spectral imaging data cubes, iteratively increasing the number of clusters from 10 to 200. For each of the 20 clustering results, cluster pixels appearing in the green area which belonged to clusters in the blue region of interest were counted to qualitatively monitor the accuracy of the clustering. This way, the number of incorrectly clustered pixels could be monitored as the number of clusters increased for each data type. For all data cubes, the number of incorrectly clustered pixels reduced as the number of clusters increased as illustrated in Fig. 2.9.

For clustering PRISMS reflectance data there were still 21621 pixels incorrectly clustered after 200 clusters were found, whereas, both colour spaces were successfully able to separate the two regions more clearly by this point. CIELAB was able to reduce the number of incorrect pixels to zero by 110 clusters and sRGB was able to reduce the total number of incorrect pixels down to as low as 64 by 200 clusters. Both clearly show a massive improvement over the use of spectral reflectance, however, CIELAB has much better performance due to its reduction of incorrect pixels to lower numbers in a smaller number of clusters than sRGB. These results imply that the best course of action when performing clustering on spectral imaging data would be to use the CIELAB data, however, there are also many other situations in which the colour of two different materials may be the same, but the spectral reflectance is different. Therefore, to maintain both the colour and spectral information it is desirable for there to be an approach which can use both CIELAB and spectral reflectance simultaneously in a complementary way.

Fig.2.8 Illustration of the two areas (in white) selected for comparing the number of incorrectly clustered pixels.

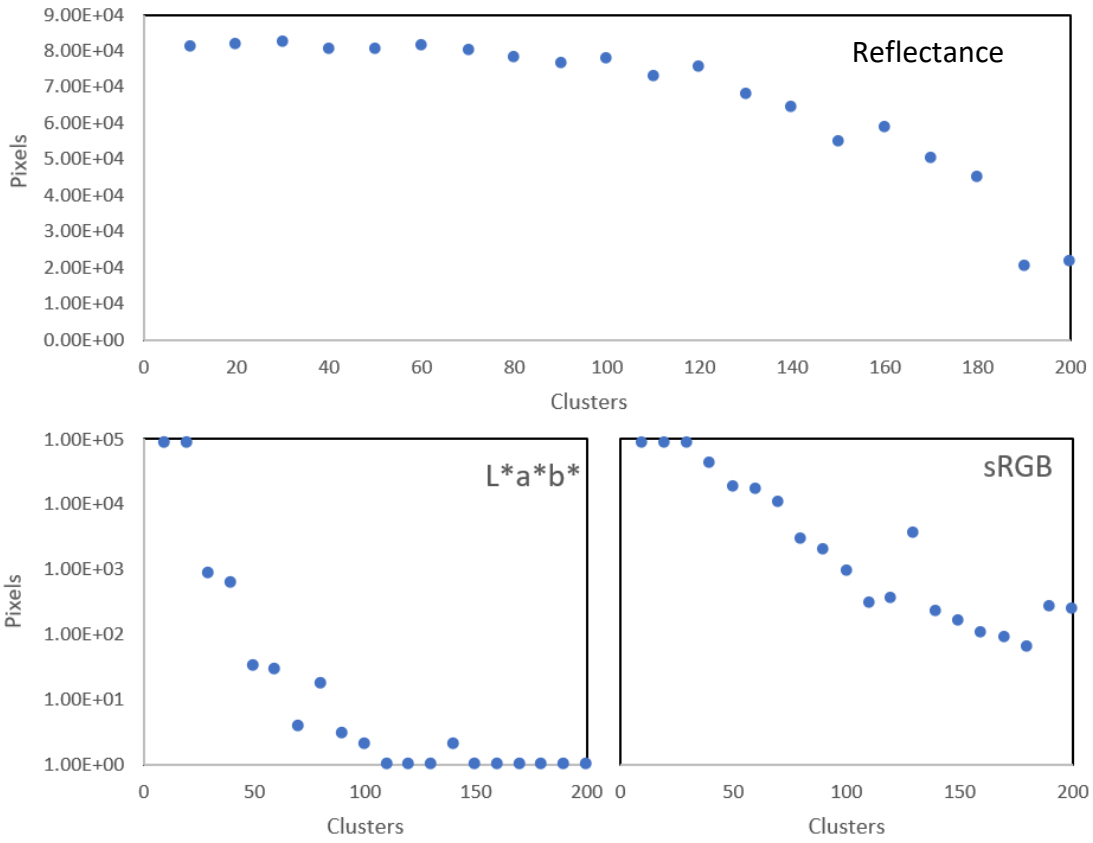

Fig.2.9 Graphs showing the decrease in the number of incorrectly clustered pixels as the number of clusters increased. CIELAB demonstrates the most accurate clustering for these tests.

2.3.3 Combined CIELAB & Spectral Reflectance Multimodal Clustering

When performing clustering for material identification at a large scale, it is important that the spectral reflectance information is not disregarded, as many materials with similar colours but clear differences in spectra may also exist, therefore, a method of performing clustering using both colour and reflectance was designed. A convenient advantage to using PRISMS spectral imaging data for this combined approach is that the colour information is derived directly from the spectral reflectance. This means that straight after processing, clustering can be performed on CIELAB and spectral reflectance together without requiring any image registration or other data treatments. Three potential methods for implementing CIELAB into the clustering process for PRISMS spectral imaging could be used: an appending approach, a map overlay approach, and a re-clustering approach. In the following sections, the map overlay, and re-clustering techniques are explored further, but a brief description of their design is given in this section.

Appending Approach

The most basic way of performing clustering on both spectral reflectance and colour data together would be to append the colour data cube onto the PRISMS spectral imaging data cube. This would generate a single 13-channel data cube where each pixel corresponds to the reflectance spectrum followed directly by the CIELAB data which can be used for clustering. However, there are problems which need to be considered when performing clustering in this way. The first challenge to consider is that each reflectance channel in the spectral data cube is effectively normalised between 0 and 1 whereas CIELAB is scaled differently and so should be normalised before performing clustering so that the differences in CIELAB do not dominate when distances between clusters and datapoints are calculated in Euclidean space. The second problem, which cannot be easily solved, is that CIELAB has only three channels whereas PRISMS reflectance has ten. The difference in channels means that clustering will be weighted more towards the reflectance, as any calculated Euclidean distances between separate datapoints and their closest SOM neurons will be dominated by differences in the spectral reflectance channels. Due to these difficulties in weighting data properly for clustering, the method of simply appending CIELAB data onto the end of the spectral reflectance was disregarded, however, in future, there may be some circumstances where such a method could be used effectively.

Map Overlay Approach

The next method of performing multimodal clustering, which will be referred to as the "Map Overlay" approach, performs unsupervised SOM two separate times to generate a pair of different reflectance and colour cluster maps where the pixel intensities correspond to the cluster number. By overlaying the two maps, a unique combination of clusters can be found which will have an equal weighting between the two approaches, avoiding the problems which would be encountered in the appending approach. A convenient example illustration of this can be seen in Fig. 2.10,

where two cluster maps are created each with two distinct clusters, which when overlayed create 4 unique clusters found across both datasets which can be illustrated on a single map.

Re-clustering Approach

The second approach, referred to as the "re-clustering" approach, would perform an initial step of unsupervised SOM on the colour information, grouping regions by their visible colour. Then, the pixels belonging to each colour-clustered group would be re-clustered at a second level by their reflectance spectrum. As a reflectance spectrum is more informative than CIELAB due to its increased number of channels (10 instead of 3) one would expect to see more finely clustered groups appear when re-clustering by reflectance within the CIELAB group. A simple graphical example of this approach can be seen in Fig. 2.11, where the 6 CIELAB groups have been found from the input data and then within the central group, defined by the green area, there are a further set of subclusters.

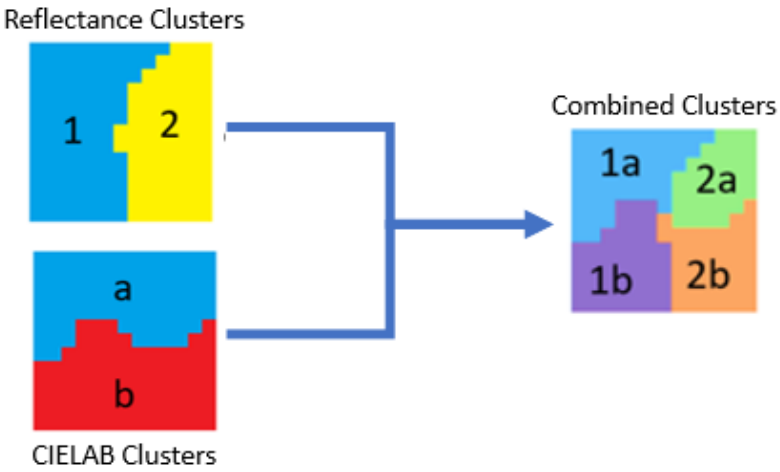

Fig.2.10 Basic diagram illustrating how map overlay clustering performs for multimodal clustering

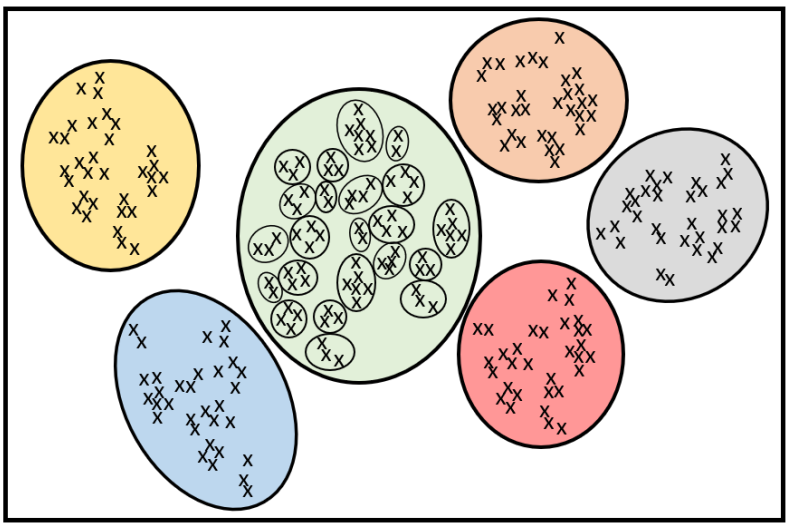

Fig.2.11 Basic diagram illustrating how re-clustering performs for multimodal clustering

2.4 SOM Based Algorithm for the Automated Clustering of Large Collections

To perform clustering on large collections, a fully scalable unsupervised clustering method that could be used for data of small to large sizes was required so that the ≈2 billion reflectance spectra in the British Library's maritime Southeast Asian collection could be reduced to a smaller number of unique spectral groups for further analysis. The designed method involved multiple stages where SOM re-clustering, SOM map overlay clustering, and a variation on hierarchical clustering, are implemented to perform the automated clustering of large collections in four steps: data reduction; repeated clustering; filtering; and merging. Though in addition there is another supplementary stage, referred to as the grouping stage, which is useful to further reduce the workload when performing analysis depending on the research questions being investigated.

2.4.1 Data Reduction

Due to the large amount of spectral reflectance data within the British Library collection, performing clustering for the entire dataset would be computationally expensive and time-consuming. To avoid such problems, the first step in performing clustering must be to involve a data reduction stage applied to each of the 1515 spectral data cubes which can group the ≈2 billion spectra into a reduced set of data. Many groups (or clusters) are expected at this stage so that each reduced spectral data cube possesses the same variation as the original data with no misclustering from requesting too small a number of clusters, however, any marked reduction in the total number of spectra will still be useful for speed improvements and automation. The idea then is that after data reduction is complete, a more complex clustering approach could then be implemented onto the reduced dataset (which will be referred to as the "reduced spectral groups") that would provide reliable and accurate clusters in a convenient amount of time for the entire collection.

For performing these operations, it can be common practise to perform dimensionality reduction techniques, such as PCA. For the clustering of spectral imaging data, techniques such as these tend to be used to reduce the number of channels in a data cube or to summarise data variation for faster processing. However, in the case of creating an automated solution to clustering spectra in large collections, this is not preferred, as precise control over specific wavelength channels or CIELAB colour bands is needed to generate accurate clusters. Furthermore, in the case of some statistical procedures such as PCA, using the technique for the entire collection does not guarantee that individual spectral cubes can be accurately compared as the principal components are dynamically computed in individual data cubes separately. It is also important to note that in a basic preliminary study, using PCA as a step prior to performing SOM did not provide any significant improvements to accuracy or speed during this reduction stage.

2.4.1.1 Comparison between re-clustering and map overlay for Data Reduction

To perform data reduction, both the re-clustering and map overlay approaches described briefly in 2.3.3 could be used, as each method results in a set of equally weighted clusters which are a product of both colour and reflectance, however, it is not clear which approach would be best at this point. To investigate this, different numbers of clusters were produced via the two implementations using the same spectral imaging test data from MSS Jav 24 f104v seen in section 2.3.3. In this investigation, clustering was performed on the data cube with the TSSE and the number of incorrectly clustered pixels being recorded as the number of clusters increased for both the re-clustering and map overlay techniques.

For the map overlay approach, the number of clusters produced for each map was steadily increased from 10 to 110, as in CIELAB, it was shown in section 3 that the minimum number of clusters required to fully separate the incorrectly clustered blue and green regions was 110. In total, the number of clusters ranged from 68 to 2358 after overlaying the CIELAB and reflectance clusters, with 2358 clusters being the position at which incorrect cluster pixels would no longer appear in the green test region of MSS Jav 24 f104v. For the re-clustering approach, the number of clusters produced ranged from 100-2100, to roughly match the same number of clusters as the map overlay approach. This involved 21 iterations where the number of CIELAB clusters asked for was increased from 10-210 in increments of 10. In each set of CIELAB clusters, the constituent cluster pixel reflectance spectra were then re-clustered into 10 groups. For example, In the final iteration, 210 CIELAB clusters were found using SOM, and then 10 spectral reflectance clusters were found within each of the 210, resulting in a total of 2100 clusters.

The changes in TSSE can be seen in Fig. 2.12, where the results for the two methods clearly show the expected steady decrease towards convergence, where acquiring extra clusters no longer significantly improved the accuracy of the clustering. Though both techniques are seemingly heading to the same position of convergence, Fig. 2.12 clearly shows that the re-clustering approach offers much greater accuracy than the map overlay approach for data reduction, where re-clustering converges at a smaller number of clusters and with a smaller TSSE. Furthermore, when monitoring how increasing the number of clusters affects the number of incorrect cluster pixels appearing in the green region of the spectral data cube, re-clustering is also shown to outperform the map overlay approach (see Fig. 2.13).

As expected, the graph shows that as the number of clusters increases, the number of incorrect cluster pixels decreases for both techniques. When comparing the two techniques, however, the re-clustering approach reaches a point of convergence far before the map overlay approach, which is expected considering the result seen for the TSSE. While the map overlay approach reaches zero incorrectly clustered pixels after 2358 clusters, the re-clustering method can reach the same position at only 1000 total reduced spectral groups.

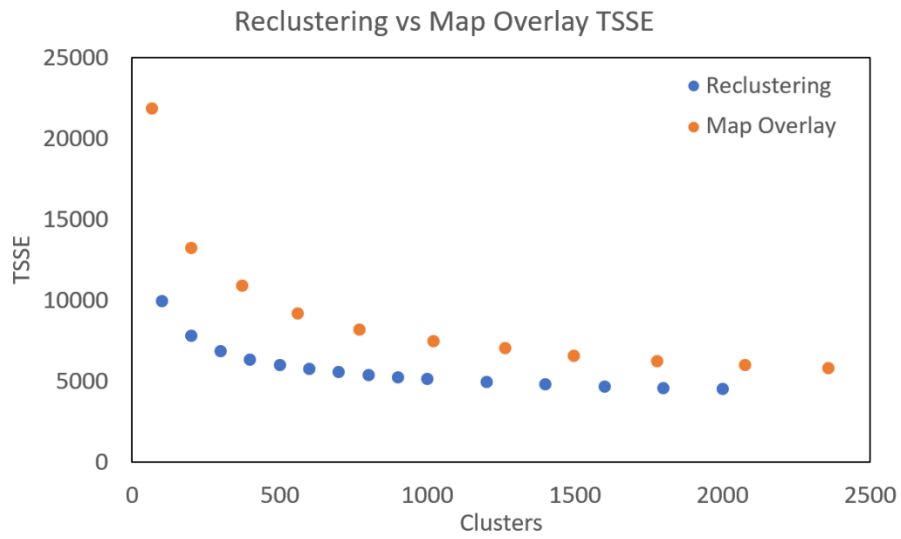

Fig.2.12 Graph showing how the TSSE changes for the map overlay and re-clustering techniques as the number of clusters increases during data reduction.

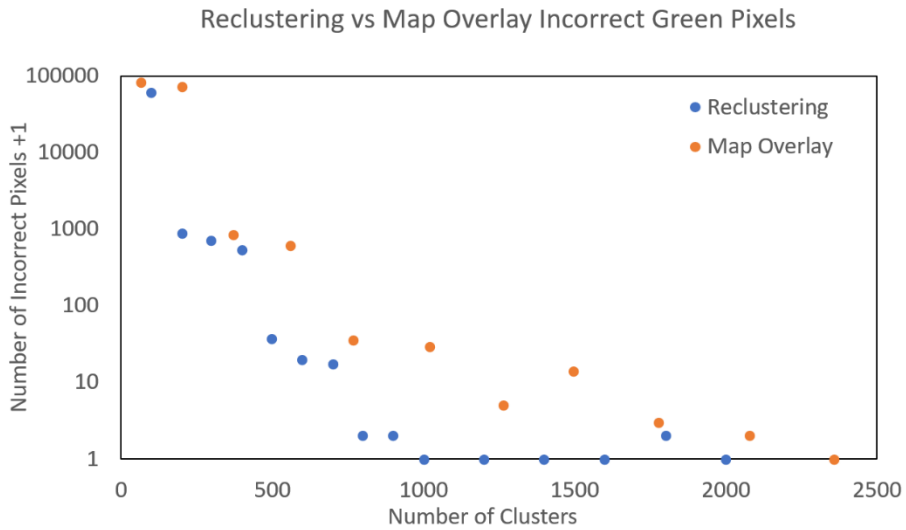

Fig.2.13 Graph showing how the number of incorrectly assigned pixels changes as the number of clusters increases during data reduction.

This is not only less than half the number of clusters required to separate the two regions in the map overlay method but also shows minor improvements as 110 CIELAB clusters were not required to perform the separation. These results suggest that the re-clustering results are also an improvement over the use of either colour information or spectral information alone, whereas the map overlay method simply combines the two with no complementary improvements above the information already available separately. It is therefore suggested by these results that re-clustering is the better approach to take when performing multimodal clustering for data reduction on spectral imaging data cubes captured by PRISMS, however in the cases of more precise clustering to be carried out at the later stages, it is not known if re-clustering will remain the most accurate approach. For this, further tests were required, as shown in 2.4.2.

2.4.1.2 Data Reduction Procedure

With the results seen in 2.4.1.1 showing re-clustering to be the better method for use directly on spectral data cubes, the final process performed for data reduction implemented this approach to acquire a set of reduced spectral groups for every data cube in the collection. This allowed for the size of the dataset for the entire collection to be reduced by a factor of 1000, which could then be used to undergo more precise clustering in a computationally efficient and realistic way. For a typical PRISMS spectral reflectance data cube, the data reduction process was found to take approximately 2-5 minutes depending on the speed of the CPU processor being used. However, for data reduction of the entire collection, the code developed can be run in parallel and in a loop to allow batches of multiple spectral data cubes to be clustered simultaneously. This allowed for significant improvements in the total time spent during data reduction and meant that large datasets can be reduced into a smaller number of reduced spectral groups in reasonable amounts of time. The process is covered in detail within the flowchart in appendix A.1.

2.4.2 Repeated Clustering

After performing data reduction, a more careful method of clustering was required to ensure that a set of precise clusters could be produced which would accurately cluster together spectra with the same reflectance and colour, whilst minimising any common misclustering mistakes such as those seen between the blue and green spectra seen previously. As this approach intends to be able to perform clustering without having priori knowledge of the complexity of the spectral data, it is difficult to gauge how many clusters should be produced for every dataset which may be encountered. Therefore, a second clustering approach was developed which transitions away from requiring user inputs such as the number of clusters, k, and replacing it instead with an approach which can automatically identify any misclustering and ensure clustering is performed correctly by using alternative parameters which hold more true meaning to spectral reflectance data. To that end, two parameters could be used to define whether a datapoint belonged to its parent cluster, and can be described as follows:

- σ_{VNIR}, a standard deviation coefficient which defines the maximum number of standard deviations about the cluster mean reflectance with which any channel of a constituent spectrum can lay before it is labelled as a misclustered spectrum.
- ΔE^*_{ab}, a colour difference threshold which defines the maximum Euclidean distance that a constituent CIELAB datapoint can have from the mean cluster CIELAB vector before being labelled as a misclustered datapoint.

The developed clustering approach uses these parameters in an iterative process in which clustering is repeatedly performed and misclustered pixels are simultaneously identified,

removed, and placed into new 'pools' to be clustered again. This process repeats until everything is placed into a cluster where all limits, set by the two input parameters, are satisfied. For example, if parameters of σ_{VNIR}=2, and ΔE^*_{ab}=10 are used, then reduced spectral groups will be repeatedly clustered until every cluster only has constituent member reflectance spectra and CIELAB vectors which lay within 2 standard deviations of the mean and have a colour difference of less than 10 respectively.

The full process can be seen given by the flowchart in appendix A.2, and each of the steps is explained in greater detail in some of the following sections of this chapter. However, prior to this, it is important to first discuss the method chosen to cluster the reduced spectral groups. While it seems reasonable to assume that the re-clustering approach should be used for the entire set of reduced spectral groups due to it performing better than the map overlay approach in the previous tests, it was not completely obvious if the results shown for a single data cube should be extended to a more complex dataset which contains information representing unique spectra spread over the entire manuscript collection. In response, the two techniques were once again investigated.

2.4.2.1 Comparison between re-clustering and map-overlay for repeated clustering.

To investigate whether the map-overlay or re-clustering approach was more accurate for performing the entire clustering process, the two methods were used to perform repeated clustering on the same reduced spectral groups followed by filtering (see 2.4.3) and then merging (see 2.4.4), as this is the procedure which would be performed to produce a final set of results for the manuscript collection. Unlike in the previous sections where the number of clusters was varied to monitor the accuracy of the techniques, an investigation was instead performed into how the TSSE changed for the reflectance data as the σ_{VNIR} changed, as this parameter serves a more 'real' purpose for setting limits on how precise the clusters should be. Therefore, for this investigation, the ΔE^*_{ab} was fixed to a value of 5, and the standard deviation was changed in steps of 0.5 from 5 down to 2, as a value of 2 was the minimum value which was used in any clustering performance on this data. This in turn produced two curves which when compared showed how the two methods affect the accuracy of the final clustering results, as seen in Fig. 2.14.

The graph in Fig. 2.14 clearly shows that for both techniques the TSSE decreases as the standard deviation coefficient becomes smaller, which is expected. However, what is interesting is that for most values of σ_{VNIR}, the map overlay approach outperforms the re-clustering method for the more complex data, which shows that re-clustering may not be useful outside of performing data reduction. One potential reason for this is likely that when performing clustering on spectrally complex data, the re-clustering approach is more prone to mistakes if the number of clusters being produced in the second level is too small.

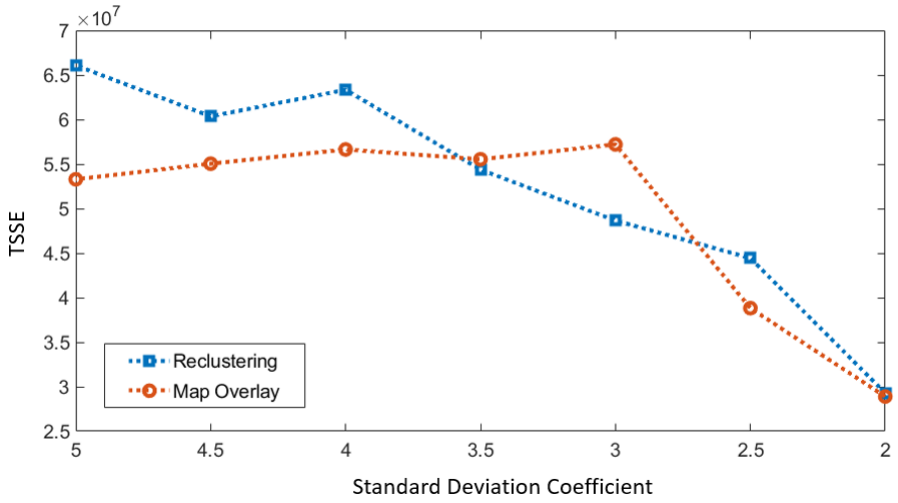

Fig.2.14 Graph showing how the TSSE for the entire collection changed as the standard deviation coefficient parameter was altered. The Map overlay method performs better than the re-clustering approach in all cases except for σ_{VNIR} = 3.5 and $\sigma_{Vis-NIR}$=3, as indicated by the lower TSSE.

This is because if a lesser number of clusters is produced by SOM than is necessary to properly separate the different spectral groups, the standard deviation for a cluster will become larger, which in turn may result in fewer misclustered spectra being removed. The map overlay method would not suffer from this problem as the two prior stages of clustering are treated independently from one another, therefore the identification of misclustered reflectance spectra would never suffer on account of the CIELAB clustering underrepresenting the true number of clusters in the collection.

While a potential fix for this problem would be easy to implement for the collection by increasing the number of reflectance clusters created at the second level of re-clustering, it is not ideal to have to adapt the number of clusters for every new dataset. This is especially true if this clustering method is to be made general enough to use across multiple collections in the future. Furthermore, there is no guarantee that increasing the number of reflectance clusters at the second level of re-clustering would marginally improve the results anyway, and it is likely that a prior understanding of the dataset would have to exist to properly optimise the re-clustering technique, which is often not the case for these large studies. Due to the results found here, the map overlay approach was used for the more precise repeated clustering, as it provides more accurate clustering for this dataset and allows for a more general methodology which can be expanded to multiple datasets.

2.4.2.2 Estimation of the Optimal Number of Clusters

Even though the method proposed for the clustering of large collections uses the σ_{VNIR} and ΔE^*_{ab} parameters during the repeated clustering stage, due to the nature of using SOM to perform the clustering on each type of data, a certain number of clusters is still required as an input. To avoid requiring priori knowledge of the data and to determine the optimal number of clusters automatically, a process was developed similar to the "Kneedle" method (Satopaa et al., 2011) to computationally estimate the elbow point of a convergence graph. To do this, the map overlay clustering method is carried out multiple times to acquire TSSE data for an increasing number of clusters ranging from 1 to $\sqrt{number\ of\ reduced\ spectral\ groups}$. In the resulting graph, a straight line is then fitted between the maximum and minimum TSSE values, and the absolute difference between the elbow plot and this straight line is computed. This produces a graph where the maximum roughly corresponds to the optimal number of clusters, which can then be used as an input for clustering (see Fig. 2.15). By performing this process prior to every repeated clustering iteration, the optimal number of clusters can be dynamically adapted to each set of original reduced spectral groups, thereby allowing for the automated accurate clustering of datasets with different sizes and spectral or colour complexity.

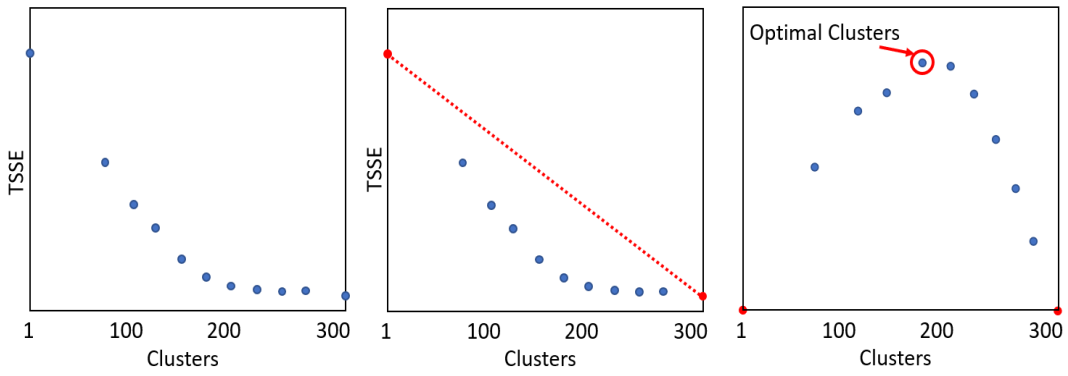

Fig.2.15 Graphs showing the process of automatically finding the optimal number of clusters. From Left: Calculate the TSSE for different cluster numbers; fit a straight line to the extremes of the TSSE plot; subtract the data from the line and find the maximum (where max = optimal clusters).

2.4.2.3 Repeated Clustering Procedure

Every time an iteration of clustering is performed, the optimal number of clusters is first estimated. Following this, the map overlay clustering on the reduced spectral groups is carried out to find a set of new 'parent clusters', which are then checked to investigate whether any misclustering has occurred within them. Identifying misclustered data is done in two parts, the first checks whether every reflectance channel for a constituent reduced spectral group spectrum lays within a certain number of standard deviations around its new parent cluster mean. The

second is performed by checking that each constituent cluster member has a ΔE^*_{ab} colour difference smaller than the user-defined parameter. Any reduced spectral groups whose reflectance or ΔE^*_{ab} lays outside the user-defined parameters are removed from their parent clusters and added to a pool to be clustered once again. Each time a reduced spectral group is removed, the parent cluster means, and standard deviation are recalculated so that the updated clusters can be checked repeatedly until there are no further misclusterings. The process is then repeated on outputted pools of previously removed groups and continues until all reduced spectral groups from the data reduction stage belong to a new cluster with no misclustered spectra.

2.4.3 Binary Cluster Filtering

From implementing colour information and reflectance spectra together, it is common for cluster maps to highlight many unique "redundant" clusters which only represent information that is not important for any kind of material identification. These often present themselves in two ways, the first is as outlines around differently coloured areas, as this is where there is normally a multiple-pixel thick boundary layer between regions which is a combination of the nearby colours or reflectance spectra. Secondly, they may appear as random single pixels due to dead or hot pixels in the CCD detector, and thirdly they can commonly occur as a result of specular reflection (commonly encountered in the collection studied in this thesis). These redundant clusters can be very high in number but are often either low in their pixel population or are very sparsely distributed random pixels throughout the image. To remove many of these redundant clusters so that the final number of clusters is a more informative and manageable set for analysis, a binary median filtering process is applied to the cluster maps directly, allowing for the removal of clusters if they only possess sparse or low population features.

The binary median filtering involves two main stages:

1. The first involves removing any clusters where the number of spectra does not exceed five in any spectral imaging data cube across the collection. For example, a cluster may have 1515 pixels but only one pixel per cube, suggesting that this is a redundant cluster.
2. Any clusters are removed if the number of cluster pixels is not greater than 5 in any spectral data cube after performing 3x3 median filtering to the binary mask which represents the distribution of cluster pixels in the 2D image (see Appendix A.3 for a flowchart).

This filtering approach ensures that the clustering results and output maps are not saturated with clusters that possess no information useful for material identification purposes. It therefore instead allows for a greater focus on regions which are more likely to hold interesting contributions for pigment identification. However, in the event that these redundant clusters may be required for specific research questions, the information is not completely lost in the process of filtering, as

any removed clusters are given a cluster label of zero and can therefore still be accessed. Filtering can be performed before or after the following merging stage and in general, should be used appropriately to remove the highest number of redundant clusters without losing important fine features. In the research for this thesis where clustering is performed on the British Library's maritime southeast Asian manuscript collection, the filtering was performed before the merging, as this offered the greatest removal of redundant clusters without losing any significant information.

2.4.4 Hierarchical Merging

The final stage, which as mentioned can be swapped with filtering depending on circumstance, involves the merging of clusters if their mean cluster reflectance spectra and CIELAB vectors match under user-defined control parameters, where for a match to exist, the CIELAB vectors of two clusters must lay within a certain ΔE^*_{ab} and simultaneously every spectral reflectance channel across both clusters must mutually lay within a certain number of standard deviations around the mean. To do this an iterative hierarchical clustering approach was designed using single-linkage criteria to allow for the fast merging of the closest matching pairs of clusters produced from the repeated SOM and or filtering stages.

This Merging process (as shown in Appendix A.4) involves:

1. Identifying the clusters which mutually overlap and producing a binary matrix where matching clusters are represented as 1, and non-matching clusters are represented as 0.
2. Producing a dissimilarity matrix computed using Euclidean distance for all clusters and normalising the matrix so that distances scale from 0 to 1.
3. Dividing the dissimilarity matrix by the binary matching matrix to produce valid normalised distances for matching cluster pairs and infinite distances for non-matching pairs.
4. Performing agglomerative clustering executed from the R stats (3.6.2) (R Core Team, 2022) package using the stats::hclust algorithm with single linkage.
5. Finding valid pairwise merging clusters in the hierarchical dendrogram and updating the mean cluster spectra, standard deviations and ΔE^*_{ab} for the new clusters.
6. Repeating steps 1-5 until no more matches exist in step 1.

Performing these stages ensures that the multiple control parameters are kept in effect during the full clustering process. Computing and combining the binary matching and dissimilarity matrices in steps 1-3 ensures that the dendrogram produced by agglomerative clustering only joins single clusters which match and possess a normalised Euclidean distance of 1 or less. To find the final clusters at step 5, the dendrogram tree is not cut by height in the traditional manner. Instead, to ensure that only matching mean cluster spectra are included in each iteration of clustering, only

the first pairwise mergings are grouped, ignoring any clustering involving hierarchically clustered groups. The result after merging will therefore consist of a set of clusters which have unique colour and reflectance which are representative of numerous individual spectra present throughout the entire collection. Each cluster will not have any potential 'copies' where two clusters with identical colour and spectra would be separated, as the merging will have placed these together, and therefore every cluster in the collection will be completely unique based both on pigment constituent, concentration, or mixture etc.

Pixel accurate calculation of statistical data

Before covering the final output data produced at the end of the clustering procedure, it is important to note that for the final merging stage, though it is used in previous stages also, calculations of mean spectra and other statistical information had to remain true to the actual clustered pixels throughout the entire collection as the data continued through multiple steps of clustering. This meant that the means or standard deviations calculated directly from all pixels in all spectral data cubes should be equal to the value produced in clustering, which is not true by default if SOM codebook vectors are used at each stage or if the mean spectra computed from the data reduction stage are carried forward to be aggregated into new means in subsequent stages. Therefore, at every step of the algorithm, every cluster mean spectrum and standard deviation had to be recalculated using statistical information stored after the stage prior.

The total sum of reflectance ($\sum$R) $\sum_{i=1}^{N} R_{i,\lambda}$, total sum of squared reflectance ($\sum R^2$) $\sum_{i=1}^{N} R_{i,\lambda}^2$, and total pixel population N, of a cluster in any reflectance or CIELAB channel can be used to perform this more accurate calculation for the mean reflectance $\bar{R}_\lambda$, mean CIELAB vector $\bar{R}_{Lab}$, and standard deviations σ_λ, σ_{Lab} using:

$$\bar{R}_\lambda = \frac{\sum_{i=1}^{N} R_{i,\lambda}}{N},\ \sigma_\lambda = \sqrt{\frac{\sum_{i=1}^{N}\left(R_{i,\lambda} - \bar{R}_\lambda\right)^2}{N}} = \sqrt{\frac{\sum_{i=1}^{N} R_{i,\lambda}^2}{N} - \left(\frac{\sum_{i=1}^{N} R_{i,\lambda}}{N}\right)^2}$$

$$\bar{R}_{Lab} = \frac{\sum_{i=1}^{N} R_{i,Lab}}{N},\ \sigma_{Lab} = \sqrt{\frac{\sum_{i=1}^{N} R_{i,Lab}^2}{N} - \left(\frac{\sum_{i=1}^{N} R_{i,Lab}}{N}\right)^2}$$

Where i is the position of the reflectance datapoint within the whole collection of data, and λ corresponds to any wavelength channel. This means that at every stage of clustering, a simple addition of the $\sum$R, $\sum R^2$, and population can be performed so that the pixel-accurate representative cluster mean spectra and standard deviations can be recalculated repeatedly without having to extract spectra directly from the original spectral cubes. This decreases computational time but

more importantly also avoids weighting issues that may arise when merging or clustering together a low-population cluster with a large-population cluster.

2.4.5 Clustering Outputs & Statistical Data

Final Cluster Maps

After filtering has been completed, final sets of statistical data for the clusters and the cluster maps are produced. The cluster maps produced are the spatial maps of the imaged object where the intensity of any pixel in greyscale value is defined by the cluster with which that pixel belongs where, for example, a pixel belonging to cluster 5 in a cluster map will have an integer intensity of 5. In addition to the greyscale maps, false colour cluster maps are also generated with unique colours representing different clusters, allowing for the convenient basic visualisation of clustering data without having to explore the cluster map integer intensities themselves.

Cluster Statistics

The statistical data which is produced after the clustering procedure is completed includes the total sum reflectance, the total sum of square reflectance, the mean cluster reflectance spectra, cluster standard deviations, and cluster pixel populations. It is necessary to continue recording this data as if any clusters eventually wish to be further merged, the quickest way to recalculate the new mean cluster spectrum and standard deviation is to use the total sum of reflectance, the total sum of squared reflectance, and the total pixel population. As mentioned at the beginning of this section, cluster grouping routines are often necessary for answering certain research questions more efficiently, therefore by storing this data in the clustering results, statistics data can be updated to provide pixel-accurate spectra for these final cluster groups.

It is important to note that at the end of the clustering process, the total number of clusters produced may still be considered very large for certain applications such as pigment identification. However, this is because the clustering code not only produces clusters with unique spectral shape but also unique intensities and colours as well. This means that many clusters for a specific pigment or mixture may exist due to factors such as changes in concentrations between certain mixture components, differences in paint layers, changes in infrared intensity due to the presence of underdrawings or differences in substrate, changes in intensity due to increases in specular reflection, and more, each of which can vary the reflectance spectra in many ways. Regardless, the reduction of billions of spectra down to thousands is still a vast improvement over performing analysis in a more traditional way, as the likelihood of missing information is removed and with the use of multiple tools and further grouping techniques, different research questions can still be answered in a much timelier manner than previously.

2.5 Automated Clustering Results for the Maritime Southeast Asian Collection

2.5.1 Clustering Evaluation for Variable-Sized Datasets

Prior to covering the different clusters created for the entire manuscript collection, a quick investigation which presented interesting results was performed to understand how changing the size and therefore complexity of the dataset affected the clustering outcome. This investigation sought to understand how the number of clusters changes with how the number of spectral data cubes introduced into the clustering method increases. It is expected for the clustering approach that as long as any new data being introduced into the method is spectrally similar, the number of clusters should eventually converge as the number of spectral data cubes increases. This is because if all the new data being introduced into the clustering method are like data already seen, they should be placed into clusters which already exist, thereby no longer increasing the total number of clusters. To investigate if this was indeed the case, the number of clusters was monitored as the number of spectral data cubes undergoing clustering was increased, the graphs for which can be seen in Fig. 2.16.

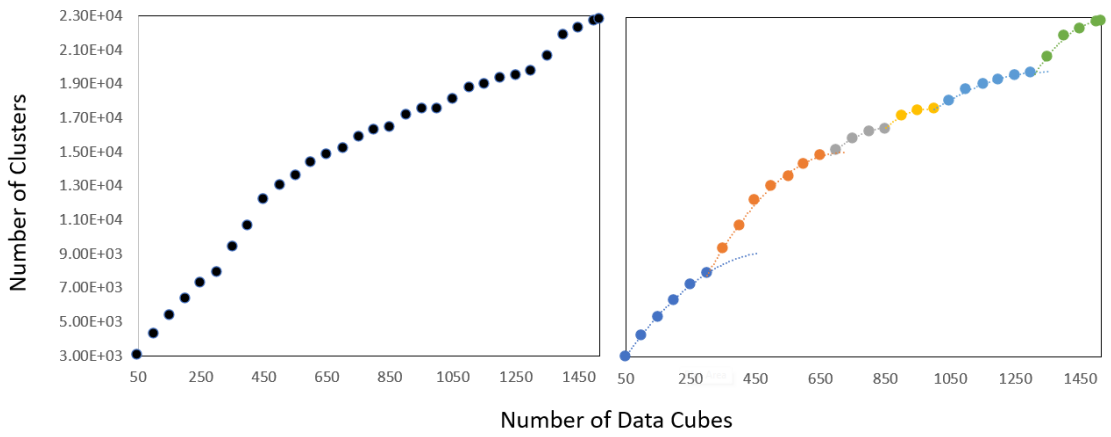

Fig.2.16 Graphs showing how the number of clusters changes as the number of data cubes being clustered increases. Note that in the graph on the right-hand side, multiple convergence curves can be found within the main trend of the entire plot.

As Fig. 2.16 shows, some level of convergence is reached in the number of clusters as the number of data cubes increases. It is interesting however that a form of convergence occurs multiple times as the number of spectral data cubes is increased. This is suspected to be because as more data is added, there are new parts of the collection which have new sets of unique spectra which have not been encountered already. This means that the position at which the graphs in Fig. 2.16 suddenly change from a trend towards convergence into a new curve, denotes the position with which a new unseen part of the collection is being clustered. As each of these smaller curves does eventually tend toward convergence, however, it seems reasonable to assume that the clustering method is performing as expected, where new spectra which are similar to previous data are increasingly added into previous clusters. It would be interesting to know with further inspection

into other datasets whether these multiple convergence curves can be used to gather a quick representation of how many artistic styles exist within the collection being clustered. However, without having access to numerous large collections this is difficult to perform, though would be interesting to understand if investigated in the future.

2.5.2 Automated Clustering Results

Using the final parameters of colour difference, ΔE^*_{ab}=2.5 and standard deviation coefficient of reflectance σ_{VNIR}= 2, a total of 22,837 clusters were found over the entire collection, an average of 15 unique clusters per spectral imaging data cube. This means that for the entire collection, the total number of spectra reaching almost two billion was reduced by a factor of ≈84200. Although the factor of reduction is significant, the overall number of clusters may still seem large at first glance, however, the reasoning for this is that many clusters do in fact have subtle differences between reflectance spectra for real reasons. Performing clustering with this level of precision allows for the separate clustering of artistic materials which can often have many subtle differences in colour or reflectance, accounting and allowing for the separation of many different properties including but not limited to differences in paint thicknesses, variations in pigment concentration, presence of underdrawings, degradation of paints, damage to paper, and spatial interfaces between materials in the 2D image plane. Some examples of these results can be seen in the following section to illustrate the strengths of performing clustering, however, some weaknesses are also discussed which cover the limitations of this clustering method when applied to VNIR reflectance and colour data.

2.5.3 Cluster Maps

As has been mentioned already, one of the main benefits of performing clustering is that the production of cluster maps can allow for a visualisation of the distribution of pigments and mixtures over many collection items at once. While most of these maps are eventually grouped at a later stage to answer specific research questions about the collection for this thesis, some of the individual cluster maps are shown in this section to demonstrate the effectiveness of performing clustering on colour and VNIR spectral imaging data. For this demonstration, multiple real colour cluster maps were produced for different manuscripts within the collection, allowing the distribution of certain commonly used pigment mixtures to be mapped, where pigments were confirmed using complementary analysis mentioned in chapter 1 and demonstrated further in chapter 3.

Mapping of Vermillion

The red pigment vermillion appears abundantly throughout the collection and is applied in many different thicknesses or with other materials depending on which manuscript is being analysed. As vermillion is used in many manuscripts, performing clustering should allow for the distribution of the pigment to be mapped over the entire collection. By creating cluster maps which show pixels for only a single cluster in their true colour, maps showing the presence of vermillion with a specific concentration and mixture can be created to aid in the analysis of the collection. A few examples of these vermillion maps can be seen in Fig. 2.17**,** where the distribution of vermillion is shown for manuscripts Add MS 12280, Add MS 12291 and MSS Jav 4.

Mapping of Green Indigo & Orpiment Mixture

Throughout the manuscript collection a green colour created by mixing indigo (blue) and orpiment (yellow), appeared in multiple different manuscripts. By applying clustering to the entire dataset, a variation of this mixture with a specific colour and pigment concentration could easily be confirmed to exist on numerous Javanese manuscripts by analysing the different cluster maps. The cluster in question, cluster 3156, was used to generate a set of true colour cluster maps which can confirm the presence of this indigo and orpiment mixture on manuscripts MSS Jav 24, Add MS 12291, and Add MS 12297. The cluster maps can be seen to demonstrate this in Fig. 2.18, where the cluster spectrum and some constituent member spectra form each spectral image can also be seen plotted together.

Mapping of Orpiment

The final example which will be shown to demonstrate the usefulness of the clustering method is for the detected yellow pigment orpiment, which can be shown to appear in many different manuscripts. Some example cluster maps showing the distribution of orpiment throughout the collection can be seen in Fig. 2.19, where cluster 527 was used to create a set of true colour cluster maps for manuscripts MSS Jav 24 and Add MS 12298.

While the clustering of orpiment performs well for most regions, there are situations in the orpiment cluster where misclustered data begins to appear. When taking a closer look at some of the pixels within cluster 527, it becomes apparent that gold can often be clustered with the orpiment due to the similarity in their colour and reflectance spectra. This situation can be seen clearly in Add MS 12291, where a closeup of the gold bulb on folio f3r (see Fig. 2.20) can illustrate both correctly clustered orpiment pixels by each of the leaf/petal motifs and misclustered gold pixels in the gold frame which have been placed together due to their similar spectra.

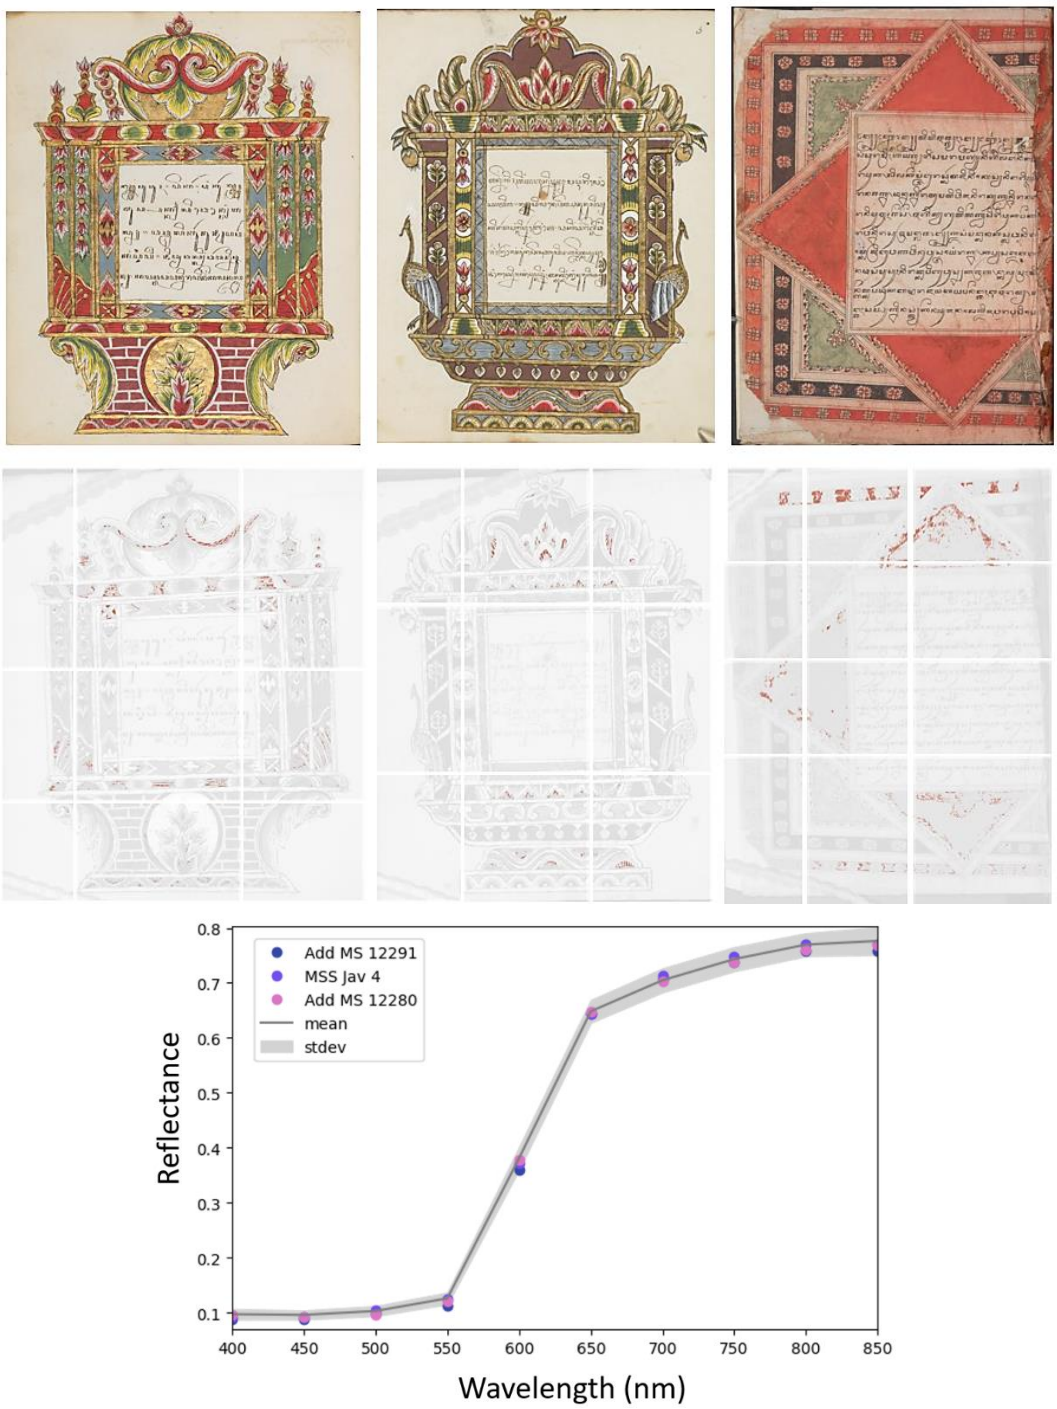

Fig.2.17 Top: RGB Images and true colour cluster maps showing the distribution of highly precise vermillion pixels over three different Javanese manuscripts. Bottom: The mean cluster spectrum of the vermillion cluster with extracted spectra from each manuscript seen fitting within one standard deviation of the mean.

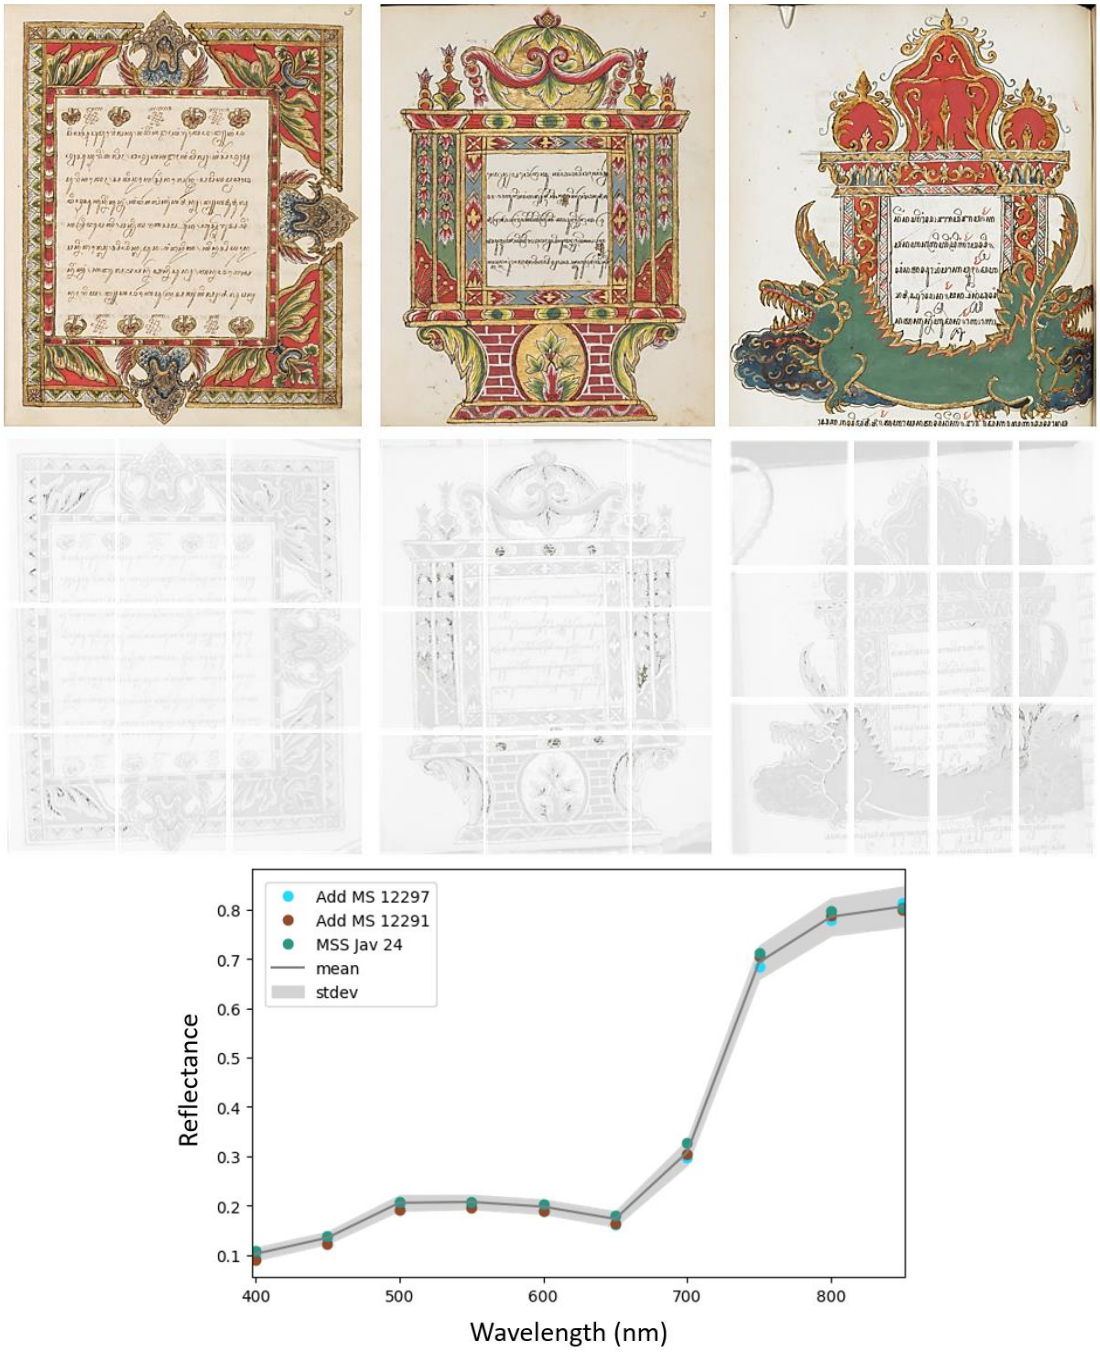

Fig.2.18 Top: RGB Images and true colour cluster maps showing the distribution of highly precise indigo and orpiment mixture pixels over three different Javanese manuscripts. Bottom: The mean cluster spectrum of the mixture cluster with extracted spectra from each manuscript seen fitting within one standard deviation of the mean.

Fig.2.19 RGB Images and true colour cluster maps showing the distribution of highly precise orpiment pixels over two different Javanese manuscripts

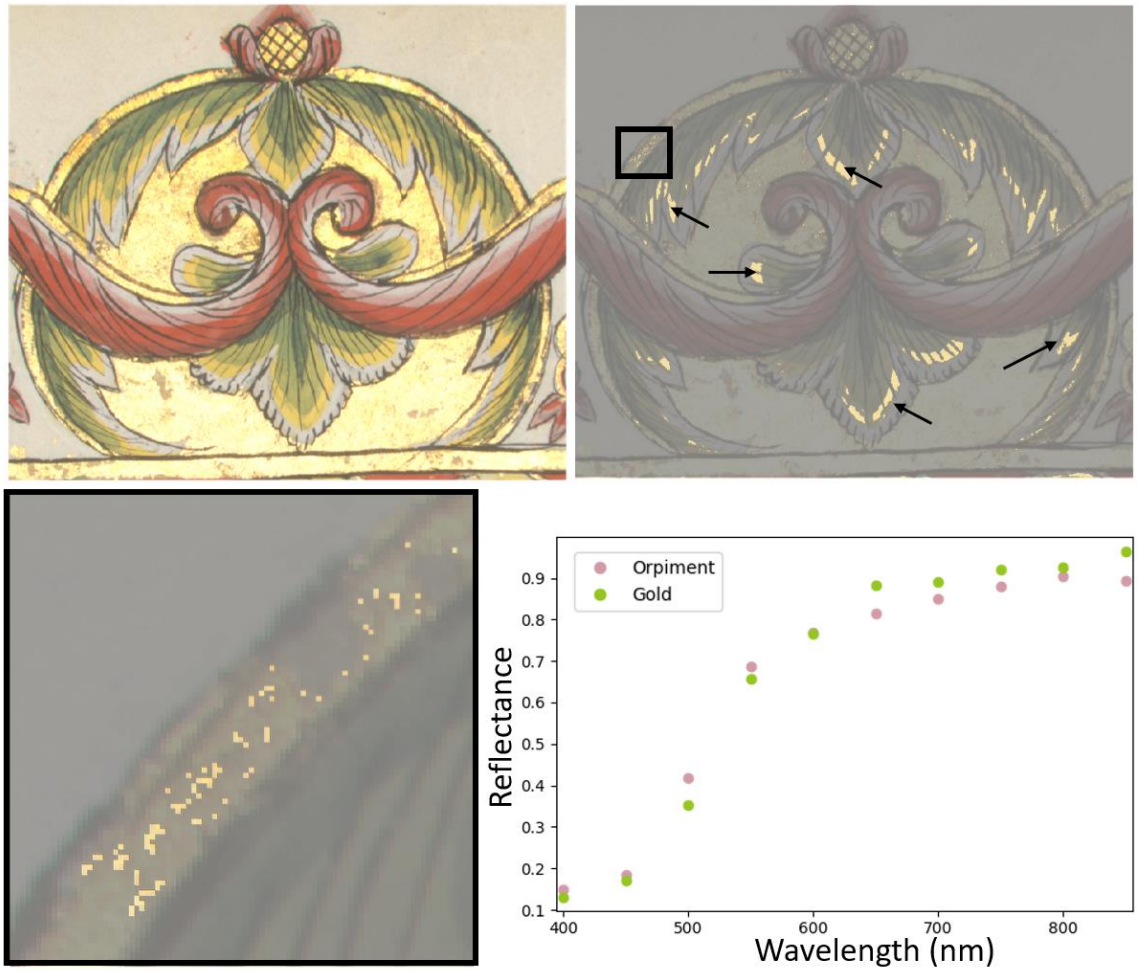

Fig.2.20 Top: RGB Image and cluster map corresponding to the orpiment cluster seen on Add MS 12291 with a region of interest shown in black for closer inspection. Bottom: Superimposed cluster map for the region of interest over the RGB image showing misclustered gold pixels which have been assigned due to their spectral similarity as shown in the graph on the bottom right. (Note: correctly assigned pixels have black arrows pointed to them and can be seen next to the green areas.)

While this is technically a scenario demonstrating misclustering, it is not a surprise that an issue like this would arise, as clustering in only one region of the electromagnetic spectrum would mean that there is no clear information which can be acquired to ensure that gold and yellow materials are always kept separate. This demonstrates a fundamental limitation of the clustering method and implies that there is no fix for this problem without introducing new data. Such a fix could be achieved through the use of alternative techniques such as XRF, which can provide the elemental composition of a material thereby allowing for gold and orpiment to be distinguished from one another easily. If XRF spectral imaging is introduced into the clustering method as an extra modality in addition to reflectance and colour, then it is likely that misclusterings such as this would not occur. This idea is explored in greater detail in chapter 4, where a holistic multimodal clustering approach is developed and performed on paper-based paintings.

2.6 Grouping of Materials and Pigment Mixtures in Large Collections

After automated clustering had been performed, it was clear that for the most part, accurate clusters can be produced which are unique to specific materials throughout multiple manuscripts, which can be useful for representing different pigments or mixtures with varying concentrations. In many research studies, however, such precise clusters are not always required, and it can be more informative to instead group together many of the clusters into fine-tuned cluster groups which can ultimately be used to answer specific questions about the data being analysed. In the case of the British Library's maritime southeast Asian collection, the focus for most of the study is to be able to map different pigment mixtures over the entire collection. As multiple different clusters can often correspond to the same mixture it would therefore be more beneficial to place these precise clusters into pigment mixture groups, which would in turn create pigment mixture cluster maps which can then be used to understand the distribution of different pigments throughout the collection.

To carry this out, a software-driven approach (see chapter 6) for the grouping of clusters was developed and used with the precise clusters formed after automated clustering of the British Library's maritime Southeast Asian collection had been performed. The full analysis of the manuscript collection can be followed in greater detail in chapter 3, however, the basic procedure for performing grouping is shown here to demonstrate a more general use of the method and to show the advantages of creating pigment mixture maps using the precise automated clustering results.

2.6.1 Grouping Procedure

The grouping procedure is a basic methodology in which clusters possessing similar mean cluster colour and spectra are manually grouped. The basic process for grouping two clusters can be described in a few steps as follows:

1. Find two similar clusters, K_A and K_B
2. For all cluster maps in the collection, convert K_A and K_B pixels with intensities I_A and I_B respectively, into a new group K_C with intensity I_C. (Where I_C is always a negative integer to ensure that original clusters and groups can be distinguished)
3. Update statistical data by combining the total sum of reflectance ($\sum R$) $\sum_{i=1}^{N} R_{i,\lambda}$, total sum of squared reflectance ($\sum R^2$) $\sum_{i=1}^{N} R_{i,\lambda}^2$, and total pixel population N. Then calculate the new K_C mean and standard deviation.

Grouping can be performed on any number of clusters at once using this approach, thereby allowing for the quick production of pigment mixture cluster maps by grouping multiple clusters with spectra which correspond to the same materials. The outputs of the grouping procedure,

therefore, include a new set of statistical data, a set of group cluster maps, and a set of false colour group cluster maps. The continued production of both the statistical data and greyscale group cluster maps allows for further grouping to be performed on already created groups at a later stage, allowing for iterative grouping to be achieved with varying levels of precision based on what the research questions are within any given study.

2.6.2 Software Implementation

When using the previously mentioned grouping technique, it was found that manually performing grouping by interacting with mean cluster spectra and colour data could prove to be very time-consuming due to the sizable number of clusters produced for large collections. In response, it was perceived that the most efficient way of disseminating results was through a software-driven approach where a researcher can visualise, analyse and group clusters together by working with all available data directly and simultaneously. This is where Guisi (see chapter 8) could be used to allow for grouping to be performed with spectra and colour through the user selection of painted areas which could be visualised in different formats as: Spectral Bands, Mean Images, Cluster Maps, True Colour Cluster Maps, sRGB Images, or False colour images. By selecting sets of clusters to group via direct interaction with spectral imaging data, subtle differences between clusters caused by variations in layering or concentrations could be investigated live. Furthermore, due to the clustering encompassing the entire collection, placing these different clusters into groups with common spectral features for only one spectral imaging data cube additionally performed grouping for the same mixtures over the entire collection.

To demonstrate an example of the grouping procedure in use on the maritime southeast Asian manuscript collection, multiple clusters belonging to the vermillion pigment discussed briefly in 5.2.1 were grouped for manuscript MSS Jav 24 on folio f104v. A colour image of this manuscript can be seen with the page highlighted in Fig. 2.21.

This involved first taking a single PRISMS reflectance spectral imaging data cube for this page and selecting multiple regions of interest (ROIs) which covered various shades of the red pigment believed to be vermillion. The 4 ROIs selected can be seen in Fig. 2.22, denoted by different coloured rectangles on the sRGB image which was created directly from the spectral reflectance data. Within these 4 ROIs, multiple clusters representing slightly different colours or spectra for vermillion could be found. When plotting the different mean cluster spectra together for each ROI, it becomes clear that the majority of clusters are largely the same, as demonstrated in Fig. 2.23, which shows the constituent mean cluster spectra within ROI 1, the leftmost region of interest in the previous image shown in Fig. 2.22.

Fig.2.21 Open-book colour image of MSS Jav 24 f104v and f105r which will be used for grouping vermillion clusters. The region which underwent PRISMS scanning can be seen highlighted by the black and yellow box.

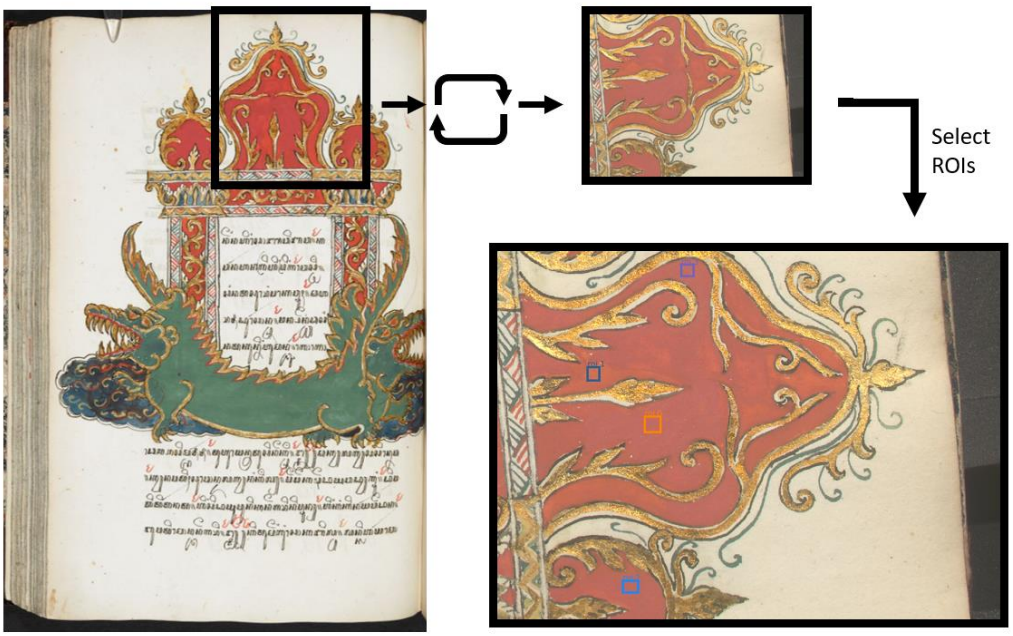

Fig.2.22 Representation of the area in MSS Jav 14 f104v which was used to perform grouping. PRISMS scanning was performed with the manuscript on its side, so the extracted region is rotated by 90 degrees for ROI selection. The ROIs used can be seen in the bottom left-hand image as small, coloured rectangles.

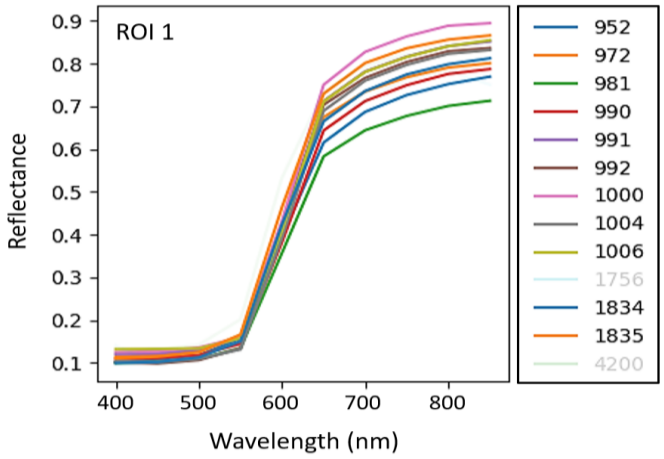

Fig.2.23 Graph extracted from the grouping software (Guisi - see chapter 6) which shows the different vermillion clusters and mean cluster spectra which will be grouped for a single ROI.

As the graph shows, most of these different clusters have similar mean cluster reflectance spectra which mainly only appear to differ in reflectance intensity towards the near infrared, implying that these areas may correspond to different thicknesses of the paint, slight concentration differences or to different shades of the substrate. Therefore, if trying to simply map where the vermillion pigment is used regardless of application, it would be appropriate to group these spectra together. Any clusters which are deemed not to match with the rest of the spectra can be ignored, this is seen for clusters 1756 and 4200, which are greyed out in Fig. 2.23 to show they will not be included in the group.

It is important to mention here that with an example such as this, a form of automated grouping can also be performed, where mean cluster spectra can be normalised and compared within their standard deviation ranges to group together spectra with similar features. However, in the case of performing clustering for the British Library's maritime Southeast Asian manuscript collection, this was not performed, and the more manual approach was taken.

When analysing all four regions of interest selected in the software, each ROI can be shown to possess its own set of unique clusters which can only be grouped if they are not already part of another ROI group. This means that for ROI 0, all constituent cluster members within the ROI can be grouped together, but for the subsequent ROI 1, only clusters which do not belong to ROI 0 can be grouped. When monitoring the unique clusters belonging to each ROI, it was shown that the final ROI possessed no constituent cluster members as all had already been grouped into the previous three. As the ROI groups also appeared to possess similar reflectance spectra, the three were combined into a single cluster group, as shown in Fig. 2.24, which would represent all the vermillion pixels in the single spectral data cube, and therefore by relation in the entire collection.

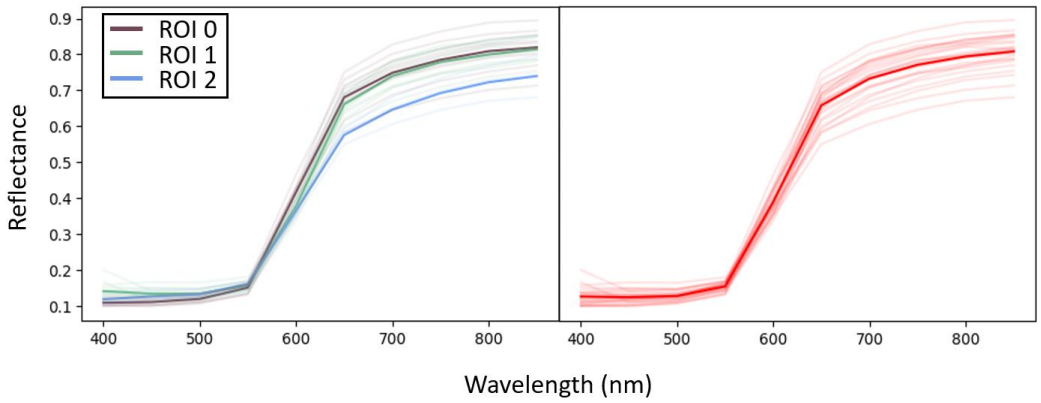

Fig.2.24 Graphs showing the constituent cluster spectra belonging to the three ROI groups (left) and then all ROIs joined into a single group given a red colour (right).

2.6.3 Grouping Results & Discussion

After performing the grouping as outlined in 2.6.1-2, pixels with spectra representing vermillion over the entire collection were all placed into the same group, after which new group cluster maps were then created. In this final section, some of these cluster maps will be shown for numerous manuscripts (see 2.6.3.1 – 2.6.3.3), but to understand how effective this method is in performing pigment mixture mapping over an entire collection, a simple table, shown in Fig. 2.25, can easily show that 35 manuscripts were found to possess the vermillion mixture grouped from only a single spectral imaging data cube.

Manuscripts with PRISMS Vermillion Group		
Add 12280	Add 12339	Mack. Pv. 42
Add 12281	Add 12363	Or 8154
Add 12284	Add 12372	Or 9333
Add 12285	Add 12379	Or 13295
Add 12287	MSS Jav 4	Or 15026
Add 12291	MSS Jav 14	Or 15227
Add 12292	MSS Jav 24	Or 15932
Add 12297	MSS Jav 28	Or 16034
Add 12298	MSS Jav 68	Or 16769
Add 12300	MSS Jav 89	Or 16915
Add 12302	MSS Malay B12	MSS Eur F148
Add 12337	MSS Malay D4	

Fig.2.25 Table showing the manuscripts shown to possess vermillion throughout the collection by grouping only the clusters in one spectral data cube.

Furthermore, as the grouping process normally only takes a few minutes to perform by having a specialist use the Guisi software, the process of mapping vermillion over many spectral data cubes using the automated clustering results within a software driven approach clearly provides huge efficiency improvements in performing reliable and accurate data analysis.

In the context of this thesis, this quick capability to map VNIR spectra over an entire collection with very little effort after clustering was a huge advantage for performing large-scale pigment identification. Furthermore, even though the groups were only created using a single data cube, large areas of grouped vermillion pixels can be identified over many different spectral cubes, instead of small areas which may happen to match well with the vermillion cluster just by chance. To demonstrate the effectiveness of the cluster grouping, three manuscripts are shown in the following subsections with their true colour “group cluster maps”.

2.6.3.1 MSS Jav 24

The best starting point to demonstrate how well the grouping procedure has performed is to show the rest of the manuscript page (MSS Jav 24 f104v) which was used to create the cluster groups in the first place. This is because if this single page of the same manuscript could not demonstrate accurately grouped vermillion pixels, then the chances of the approach working for the whole collection would be slim. When analysing the new group cluster maps created for MSS Jav 24, it became clear that the procedure had worked very well, with large regions of correctly grouped pixels being mapped not just over folio f104v, but also over other pages within the manuscript, as shown in Fig. 2.26.

Fig.2.26 RGB images and corresponding true colour group cluster maps showing the distribution of vermillion pixels assigned to the same group within MSS Jav 24.

2.6.3.2 Add MS 12291

As shown in the results directly from automated clustering, the cluster map for vermillion in Add MS 12291 demonstrated shared pixels between itself and other manuscripts. With the grouping of MSS Jav 24 being performed with a successful outcome, it is important to demonstrate whether the same performance could also be achieved on Add MS 12291, as it was known to have a provenance traceable to the same island in maritime southeast Asia (Java) as MSS Jav 24. By once again analysing the produced group cluster maps, it could be shown that manuscript Add MS 12291 had successfully undergone grouping, with maps now possessing larger areas of correctly grouped vermillion pixels, as illustrated in Fig. 2.27.

Fig.2.27 RGB image and true colour group cluster map showing the distribution of vermillion in manuscript Add MS 12291, a completely different manuscript to that used in the grouping procedure.

2.6.3.3 Or 8154

While the grouping of vermillion pixels between manuscripts with similar provenance seems logical, it was not known whether the grouping procedure would perform well with the clustering results for manuscripts from different locations in maritime southeast Asia. Therefore, a manuscript with a completely different provenance, Or 8154, originating from Bugis, South Sulawesi, was also investigated (see Fig. 2.28).

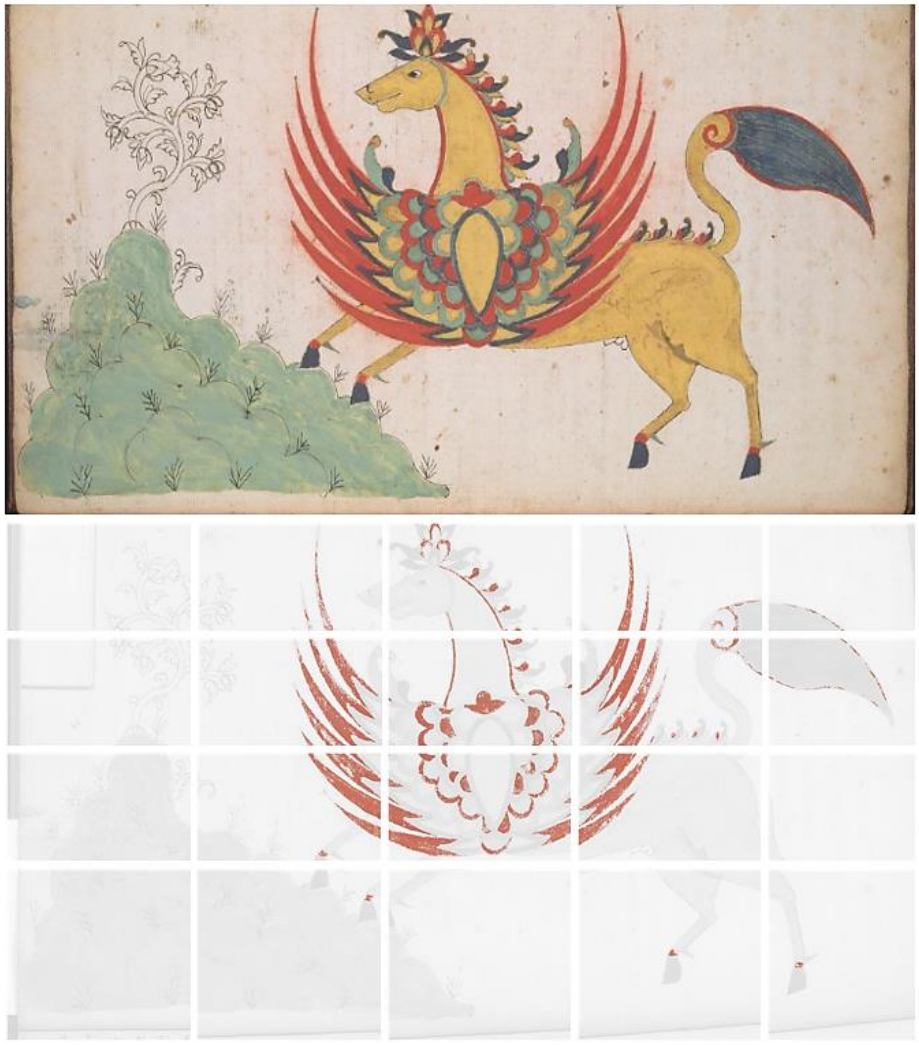

Fig.2.28 RGB image and true colour group cluster map showing the distribution of vermillion in manuscript Or 9154, a completely different manuscript to that used in the grouping procedure which also possesses a different provenance.

As can be seen from Fig. 2.28, the grouping is able to place together the vermillion clusters in a completely different manuscript with a different provenance. This therefore demonstrates that the automated clustering was able to successfully cluster vermillion of different concentrations throughout the entire collection, and with implementing this basic grouping procedure these variations can be unified if the research question only requires a mapping of basic pigment mixtures.

2.6.4 Kubelka-Munk (KM) Model Informed Grouping

With the results provided via the automated clustering method, accurate cluster grouping was shown to be accomplished by determining relationships between clusters based on their spectral shape and colour. However, this is not the only way in which grouping can be performed, in certain scenarios where the clustering of spectral imaging data is carried out simultaneously to pigment identification, much more accurate and reliable cluster grouping can be performed by using the Kubelka-Munk theory to inform a Guisi user if certain mixtures are the same. This can allow for more accurate grouping of spectra which can also account for greater differences in both spectral shape and colour, as many pigment mixtures with the same constituent materials can often vary greatly in spectral shape due to changes in relative concentration.

It was found that applying a hybrid approach of simultaneous pigment identification and clustering provided a more accurate way of grouping clusters into even fewer groups than if spectra and colour had been used alone, as greater spectral variation and colour could be confidently accounted for. Because of this, this was the process performed for the British Library's maritime Southeast Asian illuminated manuscript collection, and as a result allowed for the majority of the collection to be characterised into fewer than 200 groups in a day or so. This hybrid approach is best described by a few basic steps as shown in Fig 2.29, which demonstrates how a cluster grouping and pigment identification loop can be used to sequentially map pigment mixtures while informing where further pigment identification is required. Though a more detailed explanation is also provided in the following paragraph.

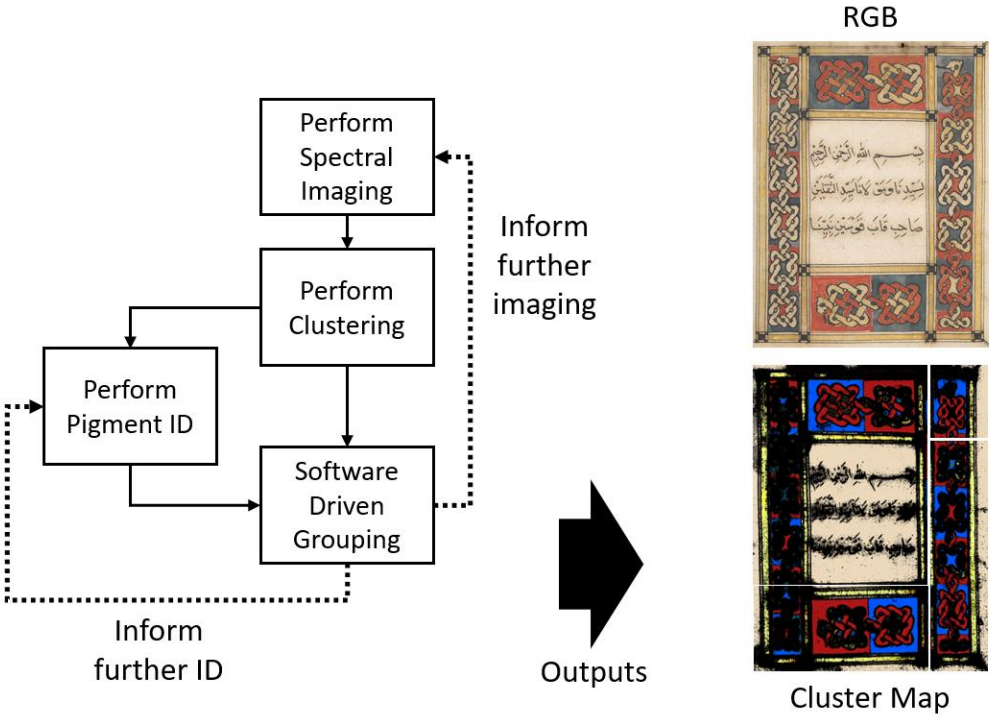

Fig.2.29 Diagram showing the procedure for KM-informed cluster grouping. An example output group cluster map created afterwards can be seen with its associated RGB image.

Essentially, the hybrid approach involves performing spectral imaging and automated clustering in the typical ways described throughout this chapter. However, before performing any grouping, pigment identification is performed so that the materials used within a small selection of mixtures are known and can be accurately modelled in the VNIR using KM mixture modelling. If a researcher is satisfied that certain VNIR spectra and colour correspond to specific mixtures, the software can then be used to group clusters which are known to represent the same pigments but in different concentrations. As this grouping will also merge similar spectra throughout the rest of the collection, every material that has been characterised by a researcher into a group can be mapped, as shown in the cluster map in Fig.2.29, which maps paper, vermillion, indigo, and orpiment in false colour. Any pixels within the cluster maps after automated clustering that have not been assigned to a group will be also highlighted in these maps as black in false colour, implying that areas with black pixels need to be further analysed or have their pigment mixture identified. Eventually, by performing grouping over many different areas where point measurements have been performed, complete reliable pigment mixture maps can be created for every identified mixture within a collection.

This approach was used to characterise most of the collection as seen in chapter 3, which meant that much of the data analysis in the VNIR could be performed with minimal effort when many different items using similar artistic materials had been imaged using PRISMS. The approach however did mean that thorough point measurements, such as with XRF, had to be performed to ensure that for the most part, mixtures with different constituent materials but similar VNIR spectra did not become incorrectly clustered or were at least could be identified as different if they did.

2.7 Conclusions

In this chapter, two different commonly implemented unsupervised clustering algorithms were compared to evaluate their performance for use in a large-scale spectral imaging study of the British Library's maritime southeast Asian illuminated manuscript collection. It was found that the most accurate approach for performing unsupervised clustering on VNIR reflectance data was the self-organising map used in its batch mode, however, it was also found that even when using this accurate clustering technique, misclustered data still occurred when using only VNIR spectra. To address these common misclustering situations, two new approaches to performing a form of multimodal clustering between reflectance and CIELAB were developed. These two techniques, termed in this thesis as re-clustering and map-overlay clustering, were used as the building blocks alongside hierarchical clustering to perform data reduction, precise repeated clustering, and merging of spectral reflectance and colour information for the large collection. The newly developed process allowed for the reduction of almost 2 billion reflectance spectra within a

collection down to only 22,837 sets of unique clusters, in which each cluster possessed its own unique mean spectrum representative of its constituent pixel spectra. As an output of the newly developed clustering method, cluster maps could be produced where every pixel in an image corresponded to a specific cluster and therefore, in the context of illuminated manuscripts, a specific material, pigment or pigment mixture with a specific relative concentration.

In many pigment identification studies, it is often not necessary to have such precise results which are offered by the automated clustering method outlined in this chapter. Therefore, another method was also developed to complement the results created after automated clustering by allowing collection-wide grouping to be performed. This was performed via a newly designed software suite known as Guisi (see chapter 6) which allowed for the direct simultaneous interrogation of clustering results and spectral data cubes, providing a method to perform targeted merging of precise clusters based on what the research questions of a study may be. It was found that by performing grouping in only a single spectral imaging data cube, the natural spectral variation which occurs in different painted areas can be exploited to group clusters over an entire collection. Furthermore, it was also found that by implementing a hybrid pigment identification and grouping approach, highly accurate pigment mixture maps could be created that not only illustrated the distribution of identified pigment mixtures in numerous manuscripts, but also provided insight into the data so that further, more accurate, analysis could be performed. It was found when using these approaches on the British Library's collection, most pigment mixtures could be characterised in less than 200 groups, with some groups still accounting for some differences in pigment concentration to avoid common misclustering issues which were encountered due to clusters possessing similar VNIR spectra, an example of which was provided in section 2.5.2, where gold and orpiment pixels were incorrectly assigned to each other.

These types of misclustering difficulties are expected to be an unavoidable problem encountered for different pigment mixtures and is likely a fundamental limitation of performing analysis in the VNIR. As both the colour and spectra can appear very similar, it is unlikely that clustering would ever be able to distinguish between these regions clearly. However, there are potential ways around this by introducing complementary mapping techniques directly into the clustering process, which is discussed further in chapter 4, and covers the holistic multimodal clustering of two export paintings.

Chapter 3

Analysis of the British Library's Collection of Illuminated Manuscripts from Southeast Asia

3.1 Introduction

3.1.1 Maritime Southeast Asian Illuminated Manuscript Collection

With the clustering and grouping methods outlined in chapter 2 successfully demonstrating their usefulness for characterising large-scale visible-near infrared (VNIR) spectral imaging surveys, the techniques could be implemented on the British Library's maritime Southeast Asian illuminated manuscript collection for the purposes of performing accurate material identification and pigment mixture mapping.

The British Library's Maritime Southeast Asian collection includes around 550 manuscripts from modern Indonesia and the Malay world. This study analysed 50 of the most important well-illuminated manuscripts of this collection which cover different genres and subjects, including Qur'ans, historical and fictional traditional stories, historical calendars, and letters, all from many islands and port towns or cities. To study this complex collection, manuscripts were separated into six well-definable regions based on their provenance, split into the following major islands and regions within Southeast Asia: Sumatra; the Malay Peninsula; Borneo; Java; Sulawesi; and Lombok. Within each region, manuscripts could be traced to locations which shared common traditions or artistic styles based on the history of the area and its trade with other regions. The main regions and origins associated with different manuscripts can be seen in Fig 3.1.

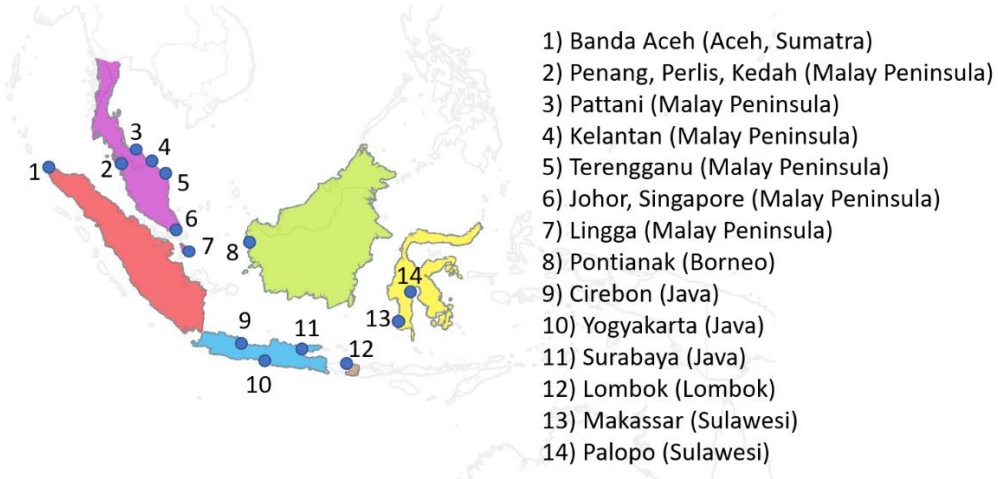

Fig.3.1 Map representing the six major regions of maritime Southeast Asia in different colours, with the individual locations of manuscript provenance seen on the right-hand side. Major regions are given in brackets.

Separating the different groups of manuscripts this way meant that comparisons between the six different locations could be carried out to investigate commonalities which may be linked by tradition, geographical similarity or through trade routes. It is also important to note however that the collection also covers a wide period, with different illuminated manuscripts being dated from as early as 1738, potentially up to the early 20th century, with the majority being from the 19th century. This implies that in addition to geographical separation, the chronological differences between manuscripts can also be used with analytical results to discover how artistic practises may have developed over time in the region.

To aid in this pursuit, a short list of every manuscript analysed along with its either known or suspected date can be seen as follows (where manuscripts with unspecified dates are marked with an asterisk:

- Add MS 12280 (1800)
- Add MS 12281 (1813)
- Add MS 12284 (1814)
- Add MS 12285 (Early 19th CE)
- Add MS 12287*
- Add MS 12291 (1813)
- Add MS 12292*
- Add MS 12297 (Early 19th CE)
- Add MS 12298 (Early 19th CE)
- Add MS 12300 (1815)
- Add MS 12302 (1802)
- Add MS 12312*
- Add MS 12337 (1812)
- Add MS 12339*
- Add MS 12363 (1759)
- Add MS 12372 (Late 18th CE)
- Add MS 12379 (Late 18th CE)
- MSS Jav 4 (1804)
- MSS Jav 14 (1814)
- MSS Jav 16 (Early 19th CE)
- MSS Jav 17 (Early 19th CE)
- MSS Jav 24 (1803)
- MSS Jav 28 (1804)
- MSS Jav 36 (1738)
- MSS Jav 67 (1812)
- MSS Jav 68 (1805)
- MSS Jav 89 (18th CE)
- MSS Malay B3 (1806)
- MSS Malay B6 (1805)
- MSS Malay B12 (1804)
- MSS Malay B14 (Early 19th CE)
- MSS Malay D3 (1806)
- MSS Malay D4 (1802)
- Or 8154 (1790-1800)
- Or 9333 (Late 19th – Early 20th CE)
- Or 13295 (1824)
- Or 15026 (1861)
- Or 15227 (Late 19th CE)
- Or 15924 (Late 19th CE)
- Or 15932 (1814)
- Or 15979 (Late 19th CE)
- Or 16034 (18th – 19th CE)
- Or 16035 (1764)
- Or 16126 (1857)
- Or 16769 (19th CE)
- Or 16915 (Early 19th CE)
- MSS Eur F148/4 F106 (1811)
- MSS Eur E378 (1813-1817)
- MSS Eur D742/1 f32r (1811)
- MSS Eur Mack Private 42 (1812)

In the following parts of this section, a brief overview of the locations is covered and includes discussions into potential subgroups of manuscripts within each region which tend to possess different artistic styles and therefore also perhaps different artistic materials.

3.1.1.1 Javanese Manuscripts

The biggest regional collection of manuscripts was those belonging to Java, of which there were 28 manuscripts with provenance from different locations such as Yogyakarta, Surabaya and Cirebon. Overall, there appeared to be four main subgroups of manuscripts within the Javanese collection which had unique artistic styles.

The first style tends to consist of a typical traditional Islamic Southeast Asian manuscript that is almost always beset with opening double frames illuminated with gold, white, red, blue, yellow, and other colours, such as green or pink, which could be created as a result of mixing those already found within the manuscript. In many of these manuscripts, the opening double frames are the only illuminations present and normally have opaque and thick paint layers, though, in some examples such as MSS Jav 24, multiple illuminated frames can be seen throughout the entire manuscript. Six examples of the first artistic style can be seen in Fig 3.2.

Fig.3.2 RGB Images of different Javanese manuscripts representing the first group of artistic styles in Java.

The second artistic style keeps a similar format as the first, with illuminated double frames always being included at the beginning of the manuscript, however, for this set, there also exists further illustration within the text. These manuscripts can include traditional Javanese Wayang-style characters, however, there are also illustrations of people in alternative styles, alongside depictions of animals, ships, buildings etc. The colour palette in these manuscripts is also like the

first style, though, through visible observation, the paints normally appear to be thinner in their application. Each manuscript includes the typical opening double frames, however, in many manuscripts, the artistic style tends to deviate slightly from the first, with the most noticeable differences being the lack of pillar illustrations and animal-based iconography. The use of gold is still included throughout and is used in the opening frames similarly to the first style. Three examples of this other artistic style can be seen in Fig 3.3, which shows MSS Jav 28, MSS Jav 68, and MSS Jav 89.

Fig.3.3 RGB Images showing numerous pages from three example manuscripts representing the second Javanese art style.

The third artistic style of Javanese manuscripts tends to have an alternative appearance, perhaps suggesting alternative influences, with preferential use of watercolour style paints and a completely different format to the previous two styles as the traditional opening double frames are not used. Furthermore, the style of illustration and figural depiction is also more westernised, with no clear use of traditional Wayang characters. Three examples of this final artistic style can be seen in Fig 3.4, showing images for manuscripts MSS Jav 67, Or 15932, and Mackenzie Private 42. All three manuscripts have very similar design structures, and the illustrations themselves are very similar also. One standout manuscript within this style is Or 15932, which also has an additional full-page painting of a brown tree over a purple and blue background, where the leaves follow a genealogical chart from Adam to Pakuwana IV and Mataram IV (Gallop, 2015).

Fig.3.4 RGB Images showing numerous pages from three manuscripts in the third Javanese art style. These possess more westernised figural depictions and illustrations.

The fourth and final artistic style for the Javanese manuscripts is unique as it does not encompass a group of manuscripts but instead covers one item within the collection: Or 9333. Or 9333 is less a single manuscript but instead appears more like a pieced-together book of numerous illuminated Wayang characters on different papers, with which many are noted to have been made by different hands. The colour palette is as extensive as some of the groups of manuscripts seen earlier, with various shades of green, blue, pink, purple, red, orange and more, applied in different thicknesses and with seemingly different amounts of degradation over time. This manuscript is the latest in the collection and is ambiguously dated to roughly the late 19th or early 20th century, therefore it is understandable why there may be differences between this and other manuscripts.

Some examples of pages from Or 9333 can be seen in Fig 3.5, demonstrating the complexity of the colours and how they are used in illuminating the different Wayang characters. This extensive use of Wayang characters marks this manuscript as being Javanese, however, it is unsure as to which port town or cities the different paintings may have originated from; however, it was suspected by Ricklefs, (1969) that they are all likely from the 19th century. If this is the case, then it would not be a surprise to see many similar pigments used withi^{n} this manuscript as may be found in some of the later manuscripts from other regions, making it another good target for analysis.

Fig.3.5 RGB images showing numerous pages within manuscript Or 9333, which encompasses the fourth artistic style of the Javanese manuscripts seen in this collection.

3.1.1.2 Peninsula Malaysian Manuscripts

Ten manuscripts originating from different locations within the Malay peninsula were also analysed from within the collection, with most possessing very similar artistic styles to those expected for most maritime southeast Asian manuscripts. Throughout this thesis, the group is commonly referred to as “Malay manuscripts” however it is important to note that this does not extend to other manuscripts which may use the Malay language, such as Sumatra, which has its own dedicated group. These Peninsula Malaysian manuscripts tend to always be illuminated with opening double frames, except for MSS EUR F148/4 f106v, and do not possess any human or animal illustrations such as the items seen from Java. There are however different styles of opening double frames which can be encountered, with the easiest way of dividing the manuscripts being those which have only red pigments, versus those which use a much wider palette. This meant that the Malay collection could be split into either a “red” and “coloured” style of manuscript volumes, and an illuminated letter. The ‘red’ set of manuscripts, uses multiple thicknesses and shades of red or pink to illuminate double frames and includes manuscripts MSS Malay B6, MSS Malay B12, MSS Malay D3 (see Fig 3.6). The manuscripts which use the wider colour palette can use colours ranging from green, to red, yellow, and blue, and include manuscripts Add MS 12379, Or 13295, Or 15227, MSS Malay B3, and MSS Malay D4, which

are shown in Fig 3.7. Finally, the letter MSS EUR F148/4 f106v has a palette which includes the use of red, black, and gold, but does not take the form of a book, and so is kept separate from the others in terms of its grouping.

*Fig.**3.6** Images representing the "red" style of manuscripts from the Malay peninsula.*

*Fig.**3.7** Images showing the "coloured" style of manuscripts from the Malay peninsula.*

3.1.1.3 Sumatran Manuscripts

Most Sumatran manuscripts in the collection are from Aceh, apart from one which is noted to be from West Sumatra. Most of the Sumatran manuscripts follow the typical Acehnese style of illumination which tends to stick with using different shades of red, normally in addition to yellow, black and reserved white. There is however one exception to this: manuscript Or 16769, which contains two sets of double frames adorned with blues and greens in addition to the traditional Acehnese palette, making it unique amongst the others. The manuscripts included in the Sumatran part of the collection include MSS Malay B14 (West Sumatra), Or 15979 (Aceh), Or 16034 (Aceh), Or 16035 (Aceh), Or 16915 (Aceh), and the previously mentioned Or 16769 (Aceh). Images showing examples of the double frames in the manuscripts can be seen in Fig 3.8, with two examples demonstrating the unique palette seen in Or 16769 (top right).

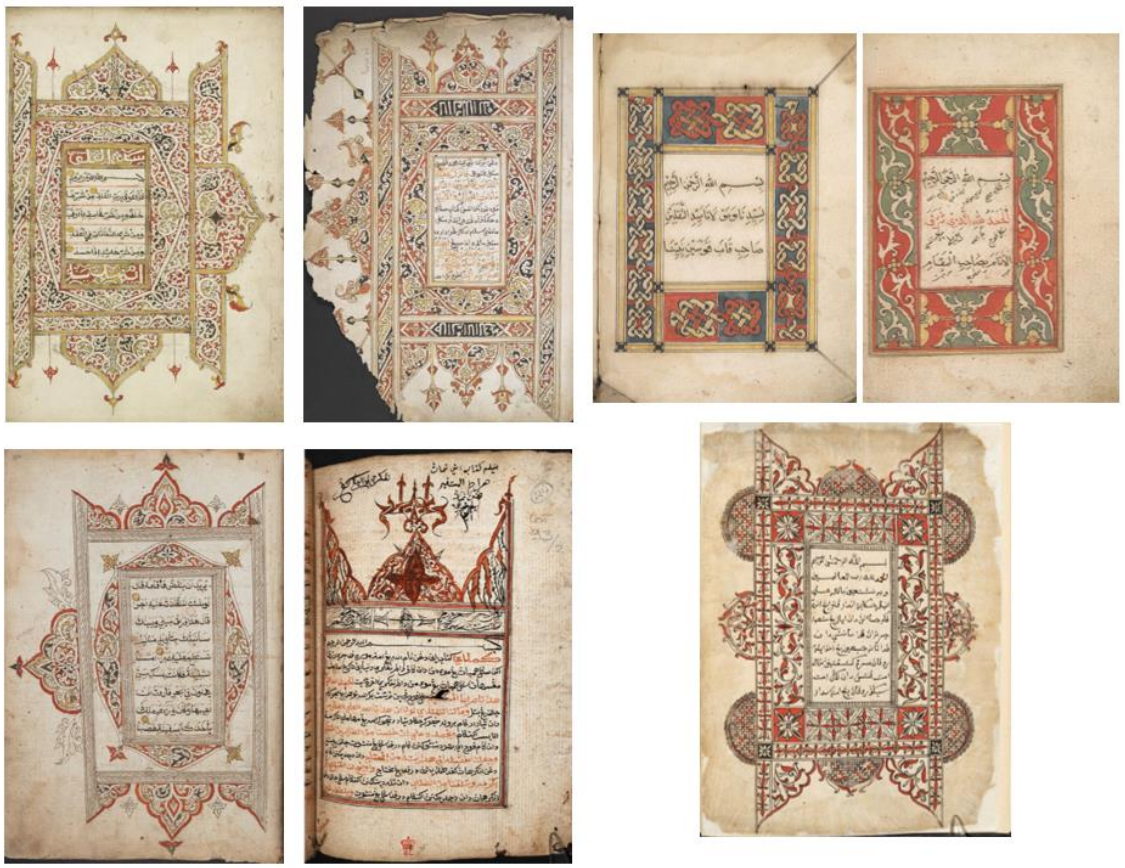

Fig.3.8 Images showing the multiple different manuscripts from Sumatra. note the standout manuscript Or 16769, which uses a different colour palette than the other manuscripts.

3.1.1.4 Sulawesi (Bugis) Manuscripts

Three manuscripts within the collection are from Sulawesi, or more specifically South Sulawesi (attributed due to their use of Bugis language and script). Each has a slightly different style, as they are not in keeping with any single traditional artistic approach, however, the palette used appears to be very similar over the three Bugis manuscripts. Images showing some illuminations for manuscripts Add 12363, Add 12372 and Or 8154 can be seen in Fig 3.9. Yellow, green and red are used on all three manuscripts, with additional use of blue colours on both Add 12372 and Or 8154 and a dark magenta colour in small roundels within the text on Add 12363 f37v.

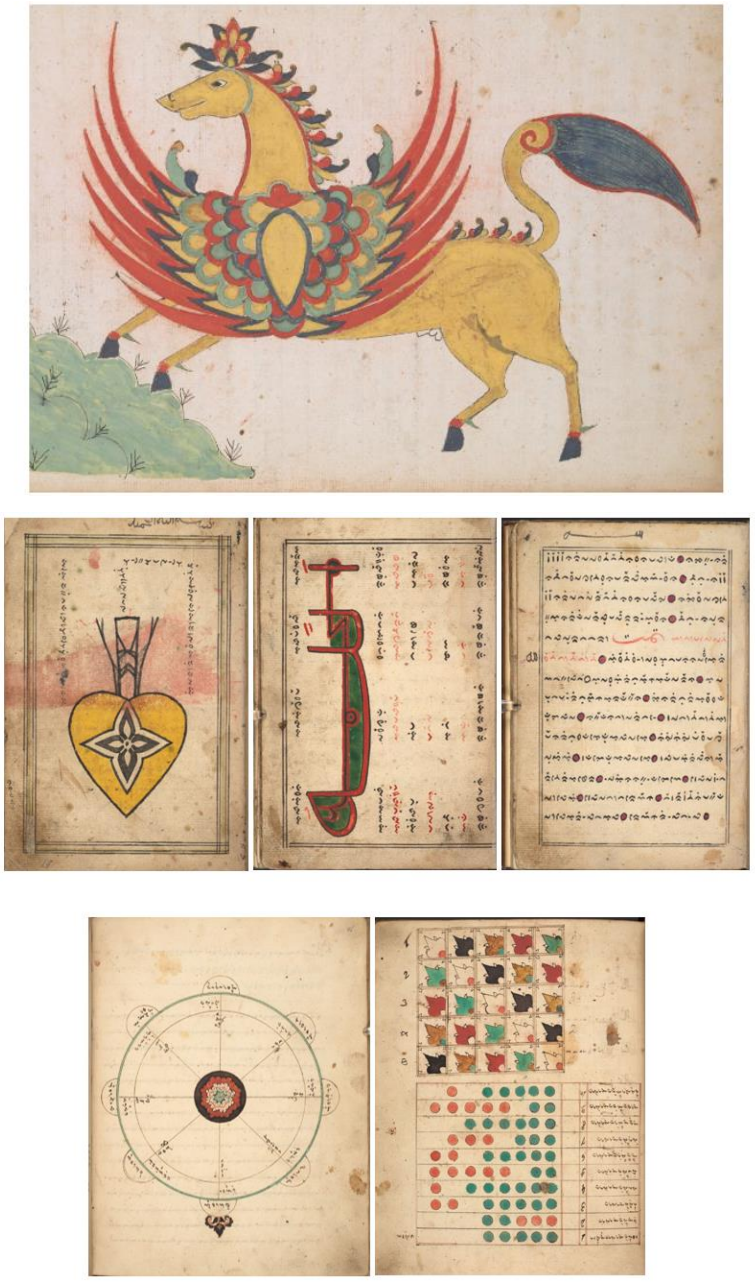

Fig.3.9 Images showing the variety of illumination in multiple Bugis manuscripts from South Sulawesi.

3.1.1.5 Manuscript from Borneo

One letter within the manuscript collection can be traced to Pontianak on the west coast of the island of Borneo. MSS EUR D742/1 f32r is a letter illuminated with a lampblack seal, dark ink and gold, an image of which can be seen in Fig 3.10.

Though not illuminated with many colours, the use of paper materials or the concentration of certain materials in the lampblack seal may give an insight into some of the potential trade connections between different islands, especially considering that most items in MSS EUR D742/1 are letters of correspondence between different Malay or Indonesian Rulers and Sir Thomas Raffles from 1810-1811, roughly the same time that many other manuscripts in the collection are also dated to.

Fig.3.10 Image showing the full letter MSS Eur D742/1 f32r, illuminated with gold, black and a lampblack seal.

3.1.1.6 Manuscript from Lombok

A single scroll most likely from Lombok with illumination on paper and a dyed silk wrap was included in the collection. Or 15924 (Fig 3.11), dated to the late 19th- early 20th century is a scroll with a sermon and prayers in Arabic but has within it some Javanese words in Arabic script. The illumination on the paper scroll is limited, with use of only black, gold, red and blue however the silk is adorned with flower-like motifs dyed red, brown, and purple. The scroll is unique within the collection, as it is the only item with original wrapping which does not consist of only paper. As this is an Arabic sermon and prayer scroll, it would be interesting to investigate whether the manuscript illumination has links with other Qur'ans or religious manuscripts such as those from Aceh or various parts of the Malay peninsula.

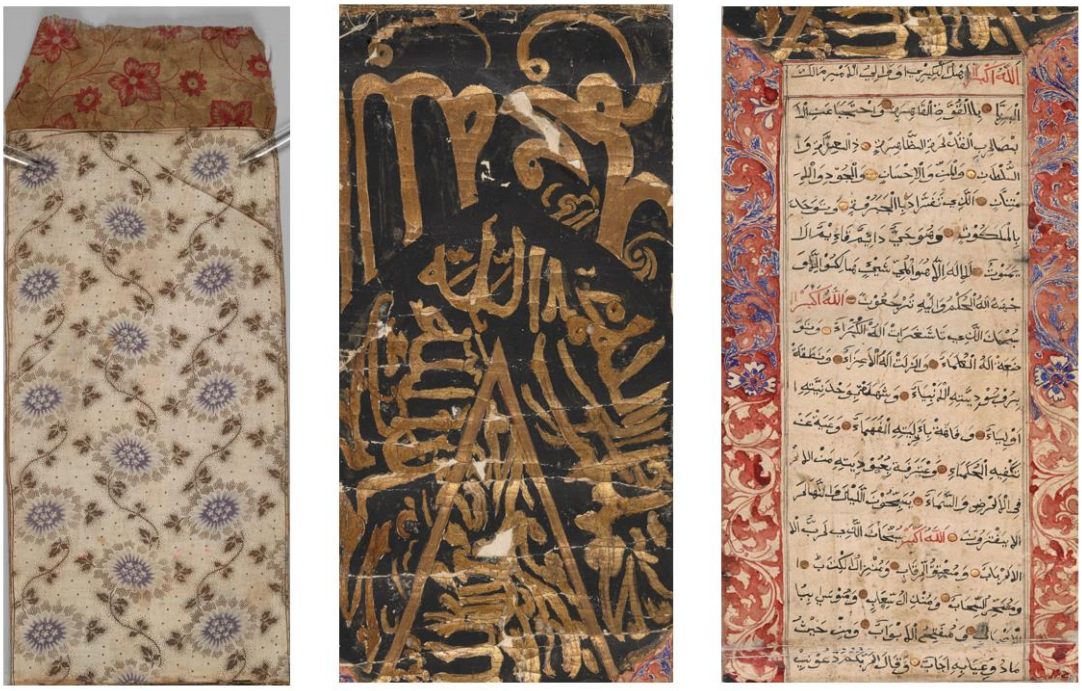

Fig.3.11 Images showing the only item within the collection likely from Lombok, Or 15924, dated to the late 19th- early 20th century.

3.1.2 Data Collection Techniques

A multitude of different scientific techniques mentioned in chapter 1 were employed to perform the large-scale spectral imaging and point measurements required for complementary pigment identification to be performed, and by extension confident Kubelka-Munk (KM) informed grouping of clusters to be achieved, as outlined in chapter 2. The different instruments used can be seen in this section and are briefly described as a quick recap so that a clearer understanding can be reached of exactly how the different methods were used together in this chapter.

3.1.2.1 Reflectance Spectral Imaging

PRISMS VNIR spectral imaging was used as the primary data collection technique to gather large amounts of VNIR reflectance information for the collection (Liang et al., 2014). For two manuscripts, however, the British Library's own system, the MegaVision EV, was used to gather this information instead (MegaVision, 2023). In both scenarios, the collected spectral imaging reflectance data could be used with other techniques in a complementary way, by extracting mean reflectance spectra from areas with which point measurements had been performed using other techniques. Furthermore, the KM informed grouping method used after automated clustering was also able to guide the complementary point spectroscopy analysis covered throughout this chapter, in turn providing large amounts of VNIR spectra that could be used holistically with other techniques to identify most artistic materials expected to be found in the collection.

PRISMS

The PRISMS data collection in this chapter is the same as described previously in chapter 2. In total for the PRISMS collected data, 1515 spectral imaging data cubes were recorded covering 48 items throughout the entire collection of 50 manuscripts, with reflectance and colour images being created for clustering and analysis purposes.

MegaVision EV

The MegaVision EV system used within this thesis offered a 14-band spectral imaging data cube ranging from 365-1040nm which could be used for analysis. It involved the use of two programmable LED arrays which could be tuned through a range of several emission bands over the entire spectral range and employed diffuser plates between the manuscript and the light source to ensure uniform illumination. The MegaVision EV system was used in the early stages of the project to record spectral imaging data for only two Javanese manuscripts: MSS Jav 36 and Add MS 12281.

3.1.2.2 UV-VNIR-SWIR Reflectance Spectroscopy

FORS – Spectral Evolution

FORS was performed using a Spectral Evolution spectrometer SR3500 (Spectral Evolution, 2023) with a halogen light source. The spectrometer covered the wavelength ranges of 350nm-2500nm, providing reflectance information from the Ultraviolet (UV) to the short-wave infrared (SWIR) with a spectral resolution of 2.8nm at 700nm, 8nm at 1500nm, and 6nm at 2100nm and a sampling spectral resolution of 1nm. The increased spectral resolution allowed for the more accurate identification of different pigments with fine absorption features in the visible range, such as cochineal or lac, and was therefore primarily used in areas where these materials were suspected. However, in a few other examples the SWIR regions were also used.

3.1.2.3 XRF Spectroscopy

XRF measurements were taken for regions with unique spectra on almost every item within the collection so that materials across different manuscripts which may possess indistinguishable VNIR spectra, but different elemental composition, could be properly identified. For the data collection of XRF throughout all the manuscripts, two separate instruments were used: the Bruker CRONO, and the Niton XL3t GOLDD+.

Bruker CRONO

The majority of XRF spectra recorded over the collection used the Bruker CRONO (Alberti et al., 2017) system fixed onto a tripod which could be placed above different manuscripts to perform point measurements non-invasively. The Bruker CRONO uses a Rhodium-target X-ray tube which functions at a variable voltage and current of 10-50kV and 5-200μA respectively, at 10W. For most measurements, a voltage and current of 40kV and 20μA were used respectively, which allowed for the identification of elements in the range of atomic number Z=14 (Silicon) to Z=92 (Uranium) without helium purge, which was the mode of use throughout this study. The CRONO can record XRF data for any given dwell time and utilises three different collimator sizes of 0.5mm, 1mm and 2mm. The 0.5mm collimator was used for most measurements as it was found to be especially useful for the analysis of manuscripts as the smaller diameter could be used for spatially precise measurements with a good signal-to-noise ratio when long dwell times (between 30-60 seconds) are used. Two guide lasers can be used with the CRONO to position the head so that the same working distance is always used when performing data collection, this allowed for the precise collection of spectra where it would be clear where the fluorescence would be induced. Furthermore, as the working distance is fixed, it avoids common issues which can arise during data collection where different count rates for elemental peaks are observed due to changes in distance between the detector and the object.

Niton XL3t GOLDD+

For a minority of the manuscripts, a Niton XL3t GOLDD+ (ThermoFisher Scientific, 2019) was used to collect XRF spectra before the Bruker CRONO was available to this study. The XL3t GOLDD+ is a handheld system which uses a silver (Ag) anode which can be used with a voltage and current of 50kV and 200μA respectively. A minimum dwell time of 30 seconds was used with a 3mm diameter spot size when recording data, which offered useful high signal-to-noise XRF spectra at the cost of spatial precision for areas smaller than 3mm.

3.1.2.4 ER-FTIR Spectroscopy

For the collection of FTIR spectra, a Bruker Alpha FTIR spectrometer was used in its external reflection mode, allowing for the non-invasive analysis of illuminated manuscripts within a range of up to 400-6000cm^{-1}. This range can allow for the identification and confirmation of many pigments and potentially some binders which may be expected to appear throughout the manuscripts. Data collection for the Bruker Alpha was also performed in-situ at the British Library by fixing the system onto a tripod, allowing for the completely non-invasive and non-destructive analysis of the illuminated manuscripts.

3.1.2.5 Raman Spectroscopy

For certain areas where the identification of pigments was otherwise difficult to perform, Raman spectroscopy was implemented. For this, a 780nm lab-built mobile Raman system was used to collect spectra so that the confirmation of certain pigments could be performed when other techniques may not be able to provide definitive material characterisation. The system used involved passing the laser through an objective to give a spot size of 15μm and could function in the spectral range of ≈65-3000cm^{-1} with a spectral resolution of ≈4cm^{-1}.

3.1.3 Analytical Techniques

In addition to data collection techniques, the analytical techniques used to identify VNIR spectra and aid in characterising the entire collection via clustering and grouping are also described.

3.1.3.1 Kubelka-Munk Model for Pigment Identification

After acquiring a VNIR reflectance spectrum either using FORS or via the extraction of spectra from spectral imaging or clustering, the Kubelka-Munk (KM) model could be used to simulate best-fitting mixtures to unknown spectra to perform pigment identification. For this purpose, several reference databases were used which contained reference spectra for different common pigments used throughout history, described in greater detail in (Kogou et al., 2015). It was often found when using the KM fitting mentioned in (Liang, 2012) that using pure references alone resulted in spectra which often did not properly fit the unknown pigment mixtures. This was normally found to be due to the contribution of the substrate on the spectra, and therefore, it was chosen to add the paper as a component during virtual mixing in nearly all circumstances, which often worked to bring the simulated mixtures closer to the unknown spectra for the suspected pigments used, in turn also improving the grouping procedures mentioned in chapter 2. While the number of spectra being mixed has no theoretical limit, it was very rare that more than 5 different constituent references were used to produce these virtual mixtures.

3.1.3.2 Automated Clustering & Grouping

In addition to the application of point analysis and prior to KM informed grouping, clustering results directly from the automated clustering method used on the entire collection of PRISMS spectral imaging data were also used to aid in the identification of paints and pigment mixtures.

Cluster-Based Analysis

Directly extracted spectra taken from PRISMS data may sometimes be difficult to use for pigment identification due to painted areas being highly absorbing or very small, and therefore possessing low signal-to-noise ratio (SNR). For these regions, the automated clustering results found in chapter 2 offer the advantage of providing spectra calculated from many more pixels than would typically be selected by using a region of interest. Additionally, automated clustering results from PRISMS could often be used to hasten the pigment identification process prior to grouping by providing their own informative cluster maps, allowing for the distribution of pigments and mixtures to be visualised more clearly. Therefore, during the analytical process, automated clustering results were also used alongside the direct extraction of VNIR spectra to allow for analysis to be performed without having to record numerous point measurements per page.

Pigment Mixture Mapping

Throughout this chapter as point measurements were recorded, the placement of clusters into groups based on their discovered pigment mixtures and or spectral shape was performed as discussed in section 2.6.3-4. This meant that differences in concentration between constituent pigments within a mixture are largely ignored, so that the same pigment mixtures can be mapped regardless of their relative pigment ratios. Furthermore, as the clustering approach developed typically ensures that severe misclusterings and misgroupings do not occur, the produced group cluster maps can normally be confidently taken as a final representative set of images which can illustrate the distribution of pigment mixtures over the collection. This therefore provides convenient visual information which can be used to aid in mapping how different islands in maritime Southeast Asia may have interacted with each other or how influences from outside the archipelago may have affected the region. In many circumstances, however, many point measurements using complementary techniques were still taken for this study across different manuscripts even if grouping had already placed painted areas in different manuscripts together, so that the performance of the clustering and grouping approach could be carefully monitored for misidentifications over the entire collection.

3.2 Material Analysis Results & Discussion

Pigment identification was primarily performed by recording point measurements throughout many different manuscripts and comparing the results against the known material composition of common pigments. Some basic pigment identification of painted areas could be easily performed by comparing different spectra against reference databases, however, for most illumination in the manuscripts, the materials used are not pure pigments and therefore do not have references. Consequently, the identification of most pigments and mixtures was regularly performed by using complementary information for different techniques, such as XRF, to perform well-informed KM mixture modelling on spectral imaging or FORS spectra. For scenarios where this approach was not enough to perform identification, further confirmation of different pigments could then be made with ER-FTIR, Raman or SWIR reflectance spectroscopy.

Throughout the collection, many different pigments were detected on manuscripts with specific artistic styles and provenance. This section will cover the pigments and mixtures detected throughout the collection, separated into parts covering the different encountered colours and likely pigments or pigment mixtures. Each subsection will mainly cover the data analysis but will also involve some discussion regarding what the results may mean for the collection when considering the provenance of the manuscripts. At the end of the chapter can then be found a final table summarising the point analysis and clustering/grouping results, followed by a collection-wide discussion/conclusion.

3.2.1 White Pigments

Throughout the collection, there is a notable lack of white pigmentation over many different manuscripts. The main reason for this is likely since in 18th-19th century maritime Southeast Asia, it was common for reserved white, where the background paper substrate is intentionally used within the illumination itself, to be used in place of white paints. Within the collection, however, there were still a handful of examples where white pigments had been applied directly onto the paper or used as a component in lightening other colours. Over the whole collection, evidence for four white pigments was found: lead white; bone ash white; barium white; and calcium carbonate-based white; all of which will be covered in the following subsections with an additional discussion into the use of reserved white, due to its common use in maritime Southeast Asian manuscript illumination.

3.2.1.1 Bone White (Bone Ash)

Bone White, or Bone Ash White, is a white pigment primarily created from bone ash, a white powder-like material produced by the calcination of bone. It can vary in form depending on the method used to produce it, however in general all types of bone white will consist of some form of calcium phosphate, implying that it can be detected through the identification of calcium (Ca) and phosphorus (P).

Javanese Manuscripts

Bone white has a history of usage in Javanese culture as a primary colourant during the preparation and painting of traditional Wayang shadow puppets, as can is described in Mellema et al., (1991), and so it stands to reason that the use of this white pigment may be expected in Javanese manuscripts, though there has been little to no evidence within any literature suggesting this to be the case for any illumination outside of traditional Wayang puppetry. Within the collection, bone ash white was detected in numerous Javanese manuscripts and was found to have been used both as a white paint applied directly onto the manuscript and in several mixtures as a whitening component to lighten other colours such as reds or blues into pinks or lighter blues respectively. There is evidence suggesting that it appears in almost every single Javanese manuscript, with it mainly presenting itself in the XRF spectra, though it was also confirmed with FTIR on multiple items. The best example to illustrate the presence of bone white is within the manuscript Add MS 12287. In this example, XRF showed increased amounts of calcium in comparison to paper and the presence of a peak corresponding to phosphorus at ≈2.01kev. With the knowledge that bone ash white mainly takes on a form of calcium phosphate, it was suspected that the constituent pigment belonging to most manuscripts with these peaks is bone ash. Though it could be further confirmed with ER-FTIR where the IR spectrum showed a strong peak at

≈1048cm^{-1} which matches well with a reference transmission spectrum of bone ash corresponding with bone ash white (CAMEO, 2020; Eastaugh et al., 2008), as seen in Fig 3.12.

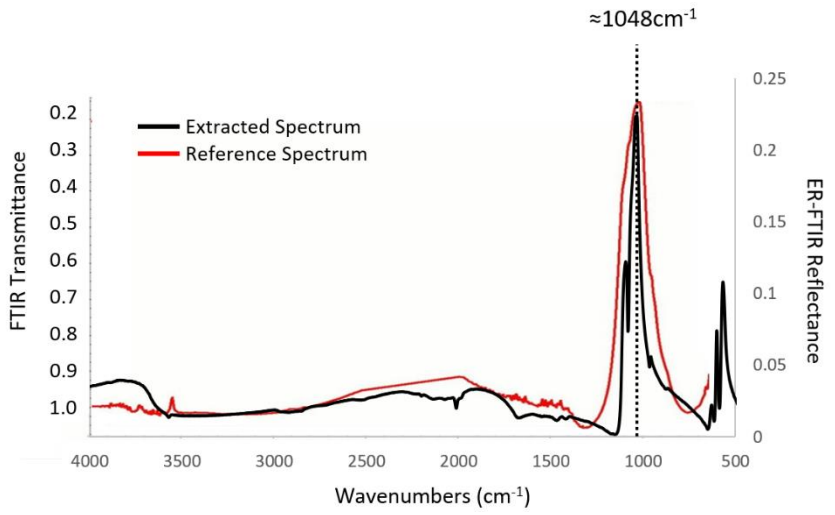

Fig.3.12 FTIR-ATR reference (Red) and ER-FTIR (Black) collected spectrum showing the similar absorption features indicative of bone ash white.

In multiple manuscripts and painted regions where this strong ER-FTIR peak existed, XRF spectra normally also possessed the Calcium (Ca) and Phosphorus (P) peaks. Therefore, it was determined that the detection of Ca and P within this manuscript collection was a good indicator for the use of bone ash white. As XRF was used in every manuscript, the detection of Ca and P could be confirmed for many other items in the collection, with at least 15 other Javanese manuscripts (see Fig 3.13) and one Malay manuscript possessing XRF peaks indicative of bone ash white. This meant that bone ash white was the most abundant white pigment used throughout the entire collection, a surprise considering that by the 18th-19th centuries, the most common white in use would have been lead white over most of the world. Furthermore, as bone ash white is known to have been used on traditional Wayang puppets, its use could suggest that some of the Javanese artists who illustrated these manuscripts also had experience in creating mixtures for and painting Wayang puppet figures. It could also be that the traditional use of colour was kept as the materials could be sourced locally.

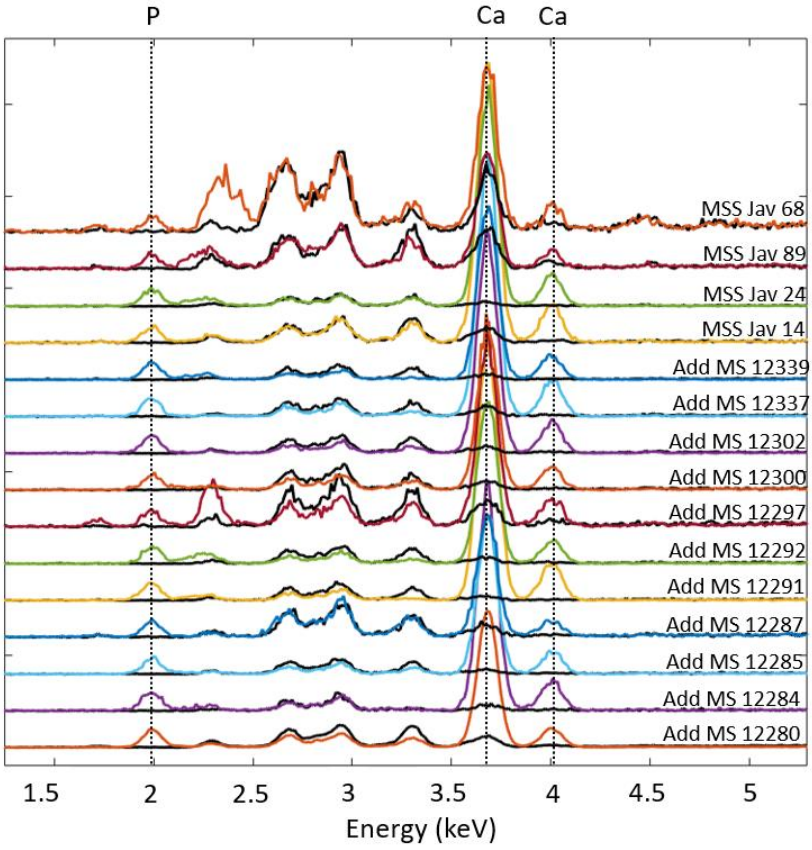

Fig.3.13 Graph showing the presence of Calcium and Phosphorus in many Javanese manuscripts. Each plot is normalised and adjusted for clarity, with the paper of the same manuscript coloured in black.

3.2.1.2 Barium White

Barium white, or barium sulphate ($BaSO_4$) could also be detected in the collection, which seems reasonable for these manuscripts as it has known to have been used as a pigment and extender in artworks from the late 18th century, commonly being a non-toxic alternative to lead white. When mixing barium white with lead white, different mixtures were often created with different ratios for different purposes, with examples including Venetian white at 1:1, Hamburg white at 2:1 and Dutch white at 3:1 Barite/Lead White (Mactaggart & Mactaggart, 2007), and so different variations may occur. It is noted in (Feller, 1986) that the minerals required to create high-quality barium white were not always available in large quantities, and that as a consequence barium sulphate was not a commonly used pigment until the early 19th century. With this being the case, it would be appropriate to assume that the pigment would appear very few times in the collection, however, if it is detected, it is likely to only appear in items which were illuminated after 1810-1820, as this was the most likely period with which it would have been made more available.

Javanese Manuscripts

Barium white was detected only once within the Javanese collection, in Or 15026 (dated to 1861), but there was evidence to suggest that it may have also been used in MSS Jav 68 (dated to 1805) and Or 9333 (late 19th-early 20th century). There is also evidence that it may have been used in other manuscripts where the primary white used was another pigment, such as lead white. To detect the pigment, ER-FTIR was used alongside XRF for the blue region of Or 15026 on the opening double frames of the manuscript. XRF peaks typical of barium were detected at approximately 4.5kev, 4.8kev, and 5.1kev, which can be seen in Fig 3.15. In addition to barium, the presence of lead could be found in the XRF spectrum, suggesting that the barium white is not pure, and is instead mixed into one of the barite/lead white mixtures mentioned previously. For further confirmation, barium sulphate could also be identified for the same region using ER-FTIR after detecting a peak at approximately 603cm^{-1}, as shown in Fig 3.16, which corresponds well with a barium sulphate transmittance reference collected using FTIR-ATR (Vahur et al., 2016) , barium white was not found in any other manuscripts in such strong concentrations but was found as a potential impurity in different lead white pigments.

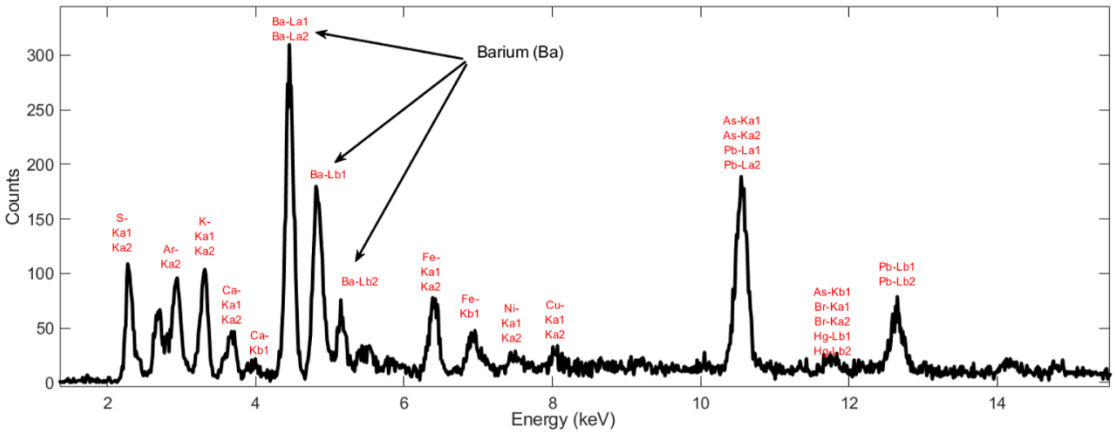

Fig.3.14 XRF spectrum showing the presence of barium along with other elements for Or 15026

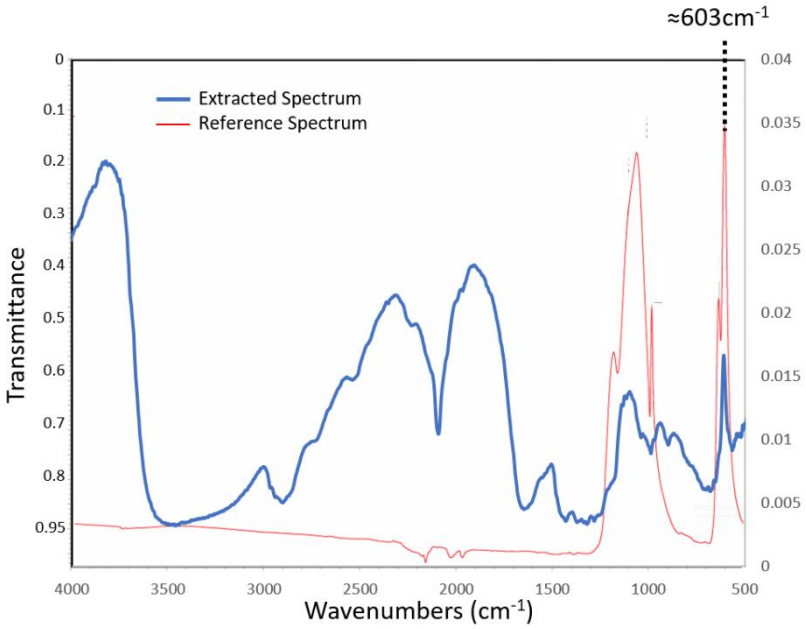

Fig.3.15 Graphs showing a (left) FTIR-ATR reference and (right) a collected ER-FTIR spectrum corresponding to barium white. Suggested by the absorption feature at approximately 603cm^{-1}.

3.2.1.3 Lead White

Lead white is considered one of the earliest pigments and has been used for many purposes, including artistic, for centuries. It is known to have been used in some of the earliest Chinese paintings, and evidence of its use in Athens, Greece can be traced to as early as 400 BC (Han et al., 2022). With lead white being used in different regions much earlier and throughout the period studied within this thesis (18th-19th centuries), it is highly likely that it would exist within maritime Southeast Asia, especially considering that both Europe and China had used lead white throughout the same centuries and had many commercial ties with the region. Within the collection, lead white was detected, however this was the case only found a handful of times, appearing in only 3 Javanese and 4 Malay manuscripts.

Javanese Manuscripts

Evidence of lead white was found in multiple Javanese manuscripts, including MSS Jav 68, MSS Jav 89, and Or 15026, in which it was always used as a component in mixtures with other pigments, as can be illustrated by analysing the normalised XRF spectra taken from different blue areas for all three manuscripts in Fig 3.17, where the main XRF lines for lead (Pb) can be seen. In each of the manuscripts however, there are additional peaks in the spectra which potentially correspond with other additional white pigments such as barium white, or bone ash white. In all three manuscripts, XRF peaks for barium can be detected, where in MSS Jav 68 f2r, this can be attributed to the presence of barium in the paper, however in Or 15026 and MSS Jav 89, the barium peaks have counts much higher than the paper, suggesting that a barium white pigment is

used alongside the lead white. In MSS Jav 68, XRF data shows the existence of phosphorus peaks along with increased counts of calcium, suggesting that lead white may have been mixed with bone ash white to further lighten the mixture.

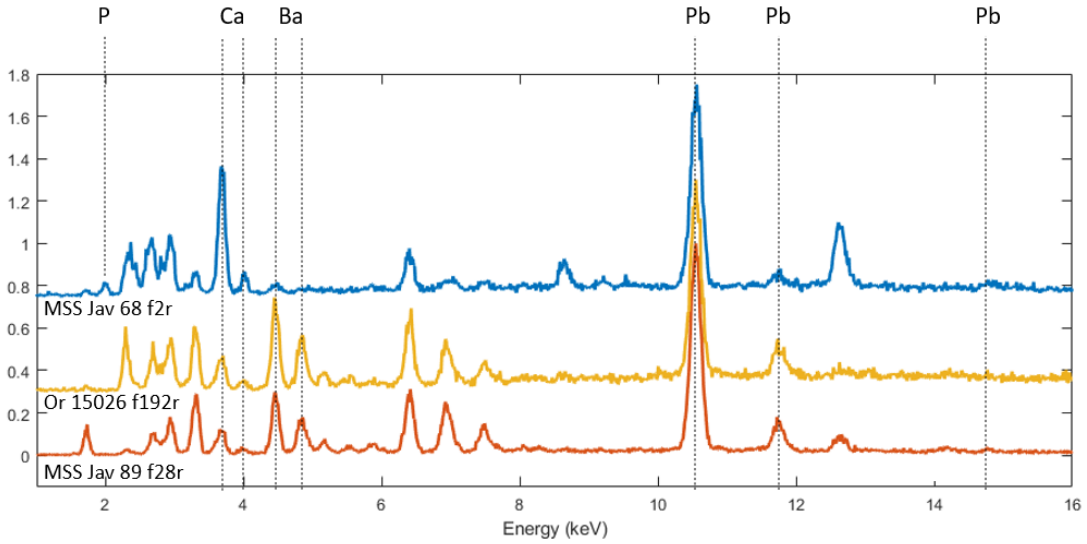

Fig.3.16 XRF spectra indicating the presence of lead for three Javanese manuscripts, indicating the use of lead white. In Or 15026 and MSS Jav 89 however, an additional presence of barium can also be detected.

Malay Manuscripts

Evidence for lead white could also be detected in MSS Malay B3, MSS Malay D4, Or 13295, and Or 15227. XRF spectra for these manuscripts can be seen in Fig 3.18, which show peaks at approximately 10.5kev, 12.6kev and 14.8kev, indicating the presence of lead white within the mixtures of different materials where other chemical elements, such as Iron (Fe), Copper (Cu), Mercury (Hg) and Sulphur (S), indicate the use of Prussian blue, copper-based greens, and vermillion respectively. For manuscripts Or 13295 and Or 15227, the use of XRF for identifying lead white can be quite definitive, as the spectra are extracted from green and blue areas which are not likely to possess other lead-based colourants. However, in MSS Malay B3 and MSS Malay D4, the areas with lead content tend to only be from regions which are red with peaks for mercury and sulphur, typically suggesting vermillion. This could instead suggest that the lead content is not present because of mixing vermillion with lead white but instead could be because of the mixing of red lead with vermillion.

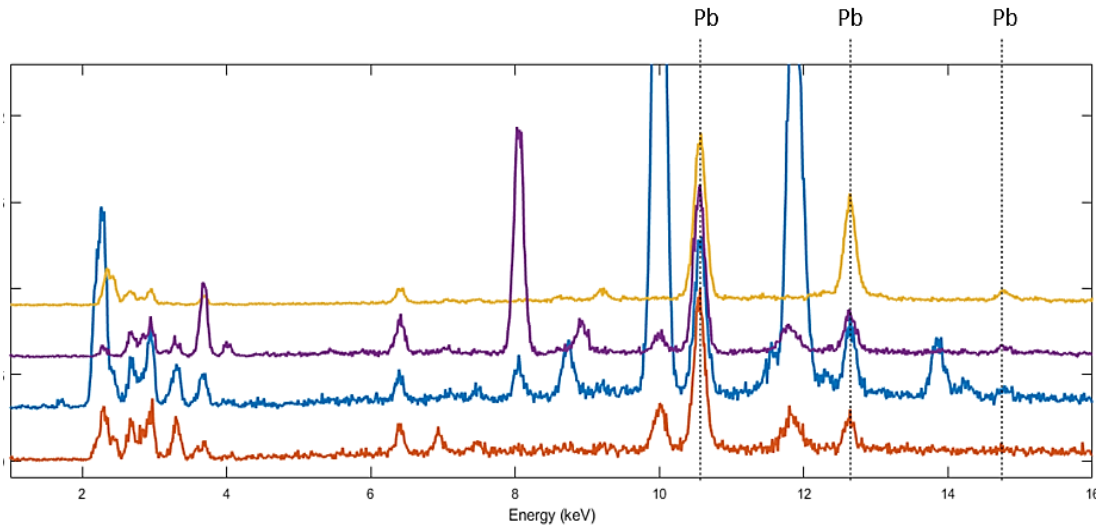

Fig.3.17 XRF spectra showing the presence of lead, indicating the use of lead white in four Malay manuscripts. From the top -- Or 13295, Or 15227, MSS Malay D4, MSS Malay B3.

3.2.1.4 Reserved White

While not a white pigment itself, reserved white has been placed in this section as it is nonetheless an artistic practice consciously used by artists depending on the school or culture to which a manuscript's illumination belongs. In Southeast Asia, reserved white is often seen in the Acehnese (Sumatra), Terengganu (Malay), and Pattani (Malay) styles (Yahya, 2021), but is not limited to these making it an interesting discussion to include if considering that there is significant use of white pigments in other manuscripts outside these locations.

Sumatran (Acehnese & West Sumatra) Manuscripts

Sumatran illuminated manuscripts are well documented to have used reserved white throughout the 18th-19th centuries, and it is often considered that its use is the most important colour within the Acehnese palette (Gallop, 2004). As expected, all six Acehnese manuscripts within the collection make use of reserved white in some form. This is even true in manuscript Or 16769, which has already been noted to deviate in artistic style from the normal examples of Acehnese opening illuminated double frames. Furthermore, in addition to these six Acehnese manuscripts, another Sumatran manuscript, MSS Malay B14 (noted to originate from West Sumatra), also uses reserved white, which is not a surprise as it possesses a style like that of Aceh or some other Malay manuscripts from Penang.

Malay Manuscripts

In addition to the previously mentioned Pattani and Terengganu schools on the Malay peninsula, some other regions have a history of using reserved white within their practice of manuscript illumination. Within the collection, many other manuscripts, whose origins include Penang, Perlis, and Kelantan, possess reserved white regions. The only exception to this could be considered to be manuscript Or 13295 (dated 31 July 1824), which consistently uses lead white both within mixtures and to decorate illuminated double frames, though this difference may be expected considering that the artistic style of this manuscript has many features more similar to those seen in Indi-Persian and Ottoman manuscript illumination (Gallop, 2013).

Sulawesi (Bugis) Manuscripts

Another island within the archipelago which normally utilises reserved white within its palette is Sulawesi, where manuscripts can have multiple parallels in artistic style similar to those in Aceh or from the Malay peninsula. A good example of one such manuscript is AKM488, a Qur'an manuscript dated to 1804 which is currently held in a collection kept by the Aga Khan Museum in Canada (Aga Khan Museum, 2023), which uses reserved white in the typical double opening frames often found throughout southeast Asia. However, a distinction should be made between Qur'anic manuscripts and those created in other regions within Sulawesi within the scope of this thesis, as the manuscripts within this collection, all created in Bugis, possess a different artistic style (as can be demonstrated in Fig 3.19). The Bugis manuscripts within the collection differ enough that they are some of the few items within the collection that do not possess the signature opening double frames so often seen in many other southeast Asian manuscripts. Furthermore, instead of being adorned with intricate patterns or motifs embedded with reserved white, the approach to implementing colour appears to be separate from many other styles. Within the British Library's collection, it could be argued that reserved white is still used for all three Bugis manuscripts, even if the artistic style used is vastly different. In manuscript Add MS 12363, a large yellow 'spade' on f40v contains an inner decorative black flower-like illustration, surrounded by empty space, or reserved white, used to accentuate the black ink. In addition, in Add MS 12372, reserved white is used in multiple flower and petal illustrations. It could also be considered that manuscript Or 8154 could be using reserved white for the tree/vine-like illustration on f3v, however, this is more likely to have just been unfinished as there are other drawings within the manuscript, such as the peacock feathers on f7v, which are not illuminated at all (see Fig 3.20).

Fig.3.18 A comparison of the different artistic styles seen in Sulawesi. The Left shows a representation of a traditional Qur'an and the right shows a Bugis manuscript illumination of Buraq.

Fig.3.19 Representation of different potential reserved white areas seen in Bugis manuscripts.

3.2.2 Blue Pigments

Throughout the collection, four blue pigments were identified using a combination of the different complementary techniques involved in this study. The most common blue pigment used in the collection was indigo, with Prussian blue being found the second most often. In addition, there is evidence for both ultramarine and azurite to have been used in the collection as well, though ultramarine only appears in two manuscripts and azurite only within one.

3.2.2.1 Indigo

Indigo is a deep blue organic pigment which has been in use over the world throughout much of history, and although its origins can be traced to many different ancient civilisations, India is largely considered to be the first major producer, processor, and distributor of indigo throughout the silk road (Balaram, 2012; Clark et al., 1993). With regions in India being some of the first civilisations to make their mark in maritime Southeast Asia (Smith, 1999; Wilkinson, 1935), it is logical to assume that indigo would have a long-standing tradition of use in many different islands within the archipelago. In its traditional form, indigo was a blue dye typically extracted from the leaves of Indigofera plants which were cultivated in abundance in Asia. Towards the end of the 19th century, the industrial mass production of synthetic indigo had become more common and the direction of trade of indigo moved from its traditional locations to those with the capacity to mass produce in factories. Many of the manuscripts in the collection were found to have contained indigo, and as most of the manuscripts range from the 17th- early 19th centuries, many are expected to contain only natural indigo dyes.

Javanese Manuscripts

Indigo is very commonly found in most Javanese illuminations and within the collection is found in every single Javanese manuscript where blue is used, though sometimes its use is not exclusive of other blue pigments. Out of the 28 Javanese manuscripts, pigment mixture maps demonstrated that indigo appeared 24 times, implying indigo to be an important aspect of manuscript illumination in Java. Within Javanese manuscripts in this collection, indigo is normally found mixed with white to create different shades of blue or can be used as a component within other mixtures to achieve different colours. A confirmation of indigo being used in multiple manuscripts can be seen in Fig 3.21, where Kubelka-Munk fitting has been used to create virtual mixtures that fit well to numerous shades of indigo. While indigo is very common, its implementation can vary depending on the manuscript being analysed, where different indigo mixtures can be seen to be used throughout the Javanese part of the collection. In the Javanese manuscripts the main indigo mixture used tends to include bone ash white, and can be seen in multiple manuscripts, for which XRF and VNIR results are illustrated in Fig 3.22, showing the

presence of typical indigo features in the PRISMS spectra and both Phosphorus and Calcium (the main components of bone ash) in the XRF.

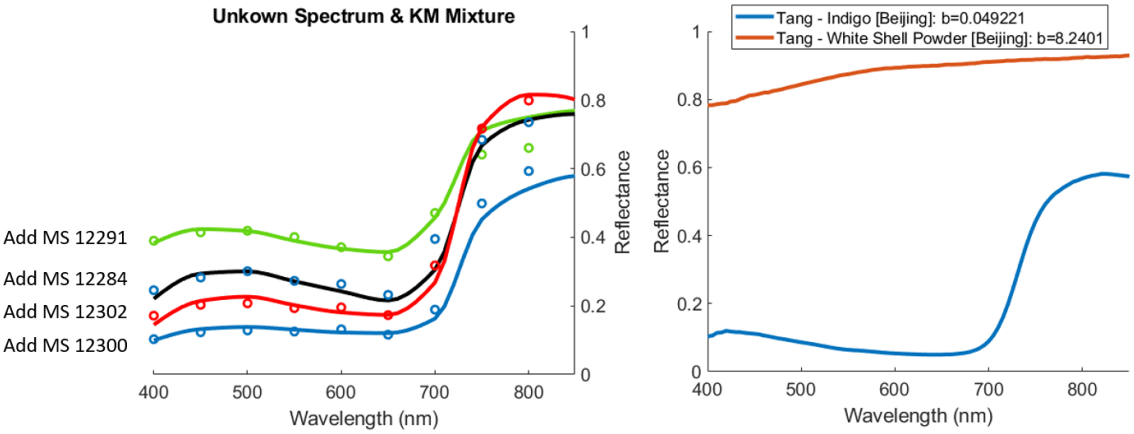

Fig.3.20 Graph showing (Left) various Kubelka-Munk fits (solid lines) for different blue indigo spectra (circles) regions in four Javanese manuscripts.

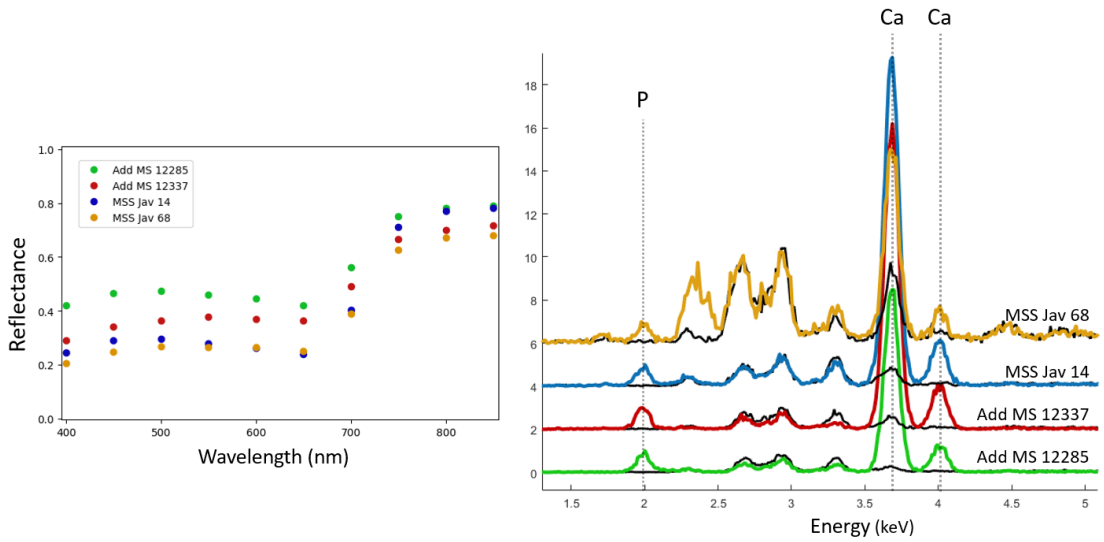

Fig.3.21 Graphs showing the presence of indigo and bone ash white mixtures in four Javanese manuscripts.

Malay Manuscripts

Indigo also appears in manuscripts Add MS 12379 and MSS Malay B3, both noted to be of Malay origin (Fig 3.23). In Add MS 12379, indigo is used to illuminate the opening double frames with a light blue colour, achieved by mixing indigo with the calcium-carbonate-based white. In MSS Malay B3 it is used for much darker regions which depict flower petals located on folios f36v and f37r, and there is evidence to suggest that it is also mixed with a yellow to create the green as well. Clustering results prior to grouping also showed the potential presence of indigo in manuscript Or 15227, a Malay Quran which typically does not use indigo throughout, where a faded drop of blue pigment can be found on f302v which has the VNIR spectral features indicative

of indigo. Whether this is a mistake made by the original artist, or if it was an accidental drop afterwards is not known, however, if it was placed during the manuscript's creation, its presence does indicate that somebody with access to indigo as a refined pigment actively chose not to use it in the illumination of the manuscript.

Fig.3.22 Images showing indigo regions in Malay manuscripts Add MS 12379 and MSS Malay B3

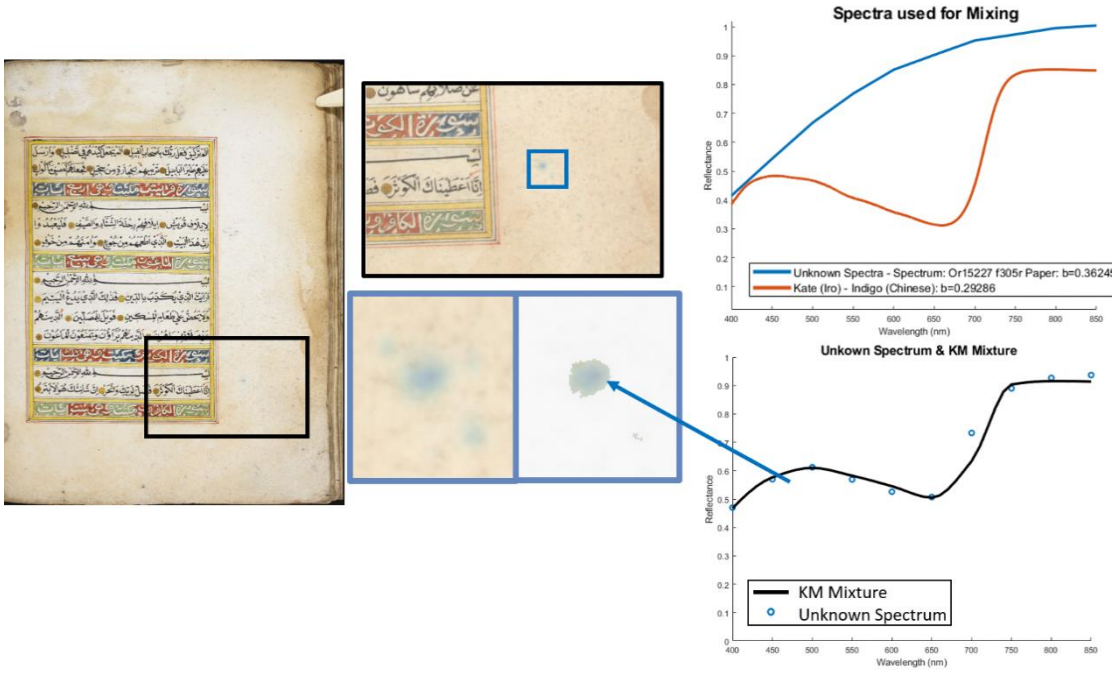

Fig.3.23 Images and graph showing the how clustering detected the presence of a drop of indigo on Or 15227, which typically uses Prussian blue throughout.

Sumatran (Acehnese) Manuscripts

Acehnese manuscripts typically do not include blue within their palette, however, within the collection there is one example which implements colours outside of the normal practice. Or 16769 contains indigo on folios f.2v and f.3r, an example of which can be seen in Fig 3.25. This

use of colour outside the red, yellow, black, and reserved white typically seen is very unusual and could suggest an influence from outside of Aceh. It is difficult to determine which culture, but some guesses can be made where, for example, analysing the XRF spectra for the blue regions in Or 16769 suggests that there is no additional white pigment added to the indigo (see Fig 3.26) making it unlikely that there would have been Javanese or Malay influence as the use of bone ash white or lead white/barite would be expected respectively. Instead, the only additional materials seen are small amounts of copper (Cu) and Arsenic (As), both minerals which can sometimes appear as impurities in natural indigo which result from its production (Qi-yue et al., 2020).

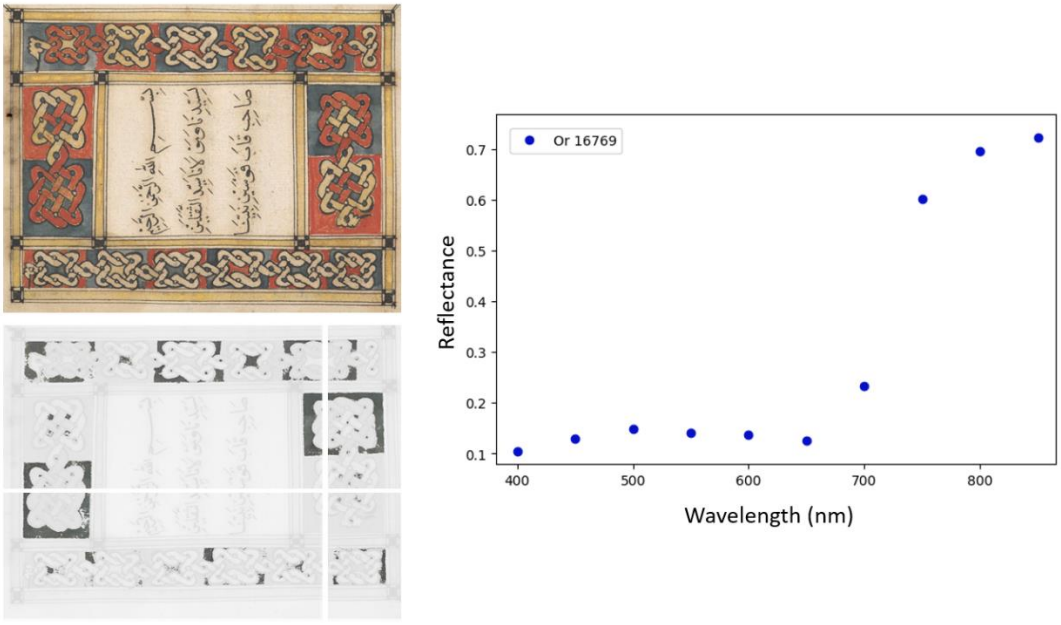

Fig.3.24 Graph showing a VNIR spectrum indicative of the indigo cluster for the blue regions in Or 16769

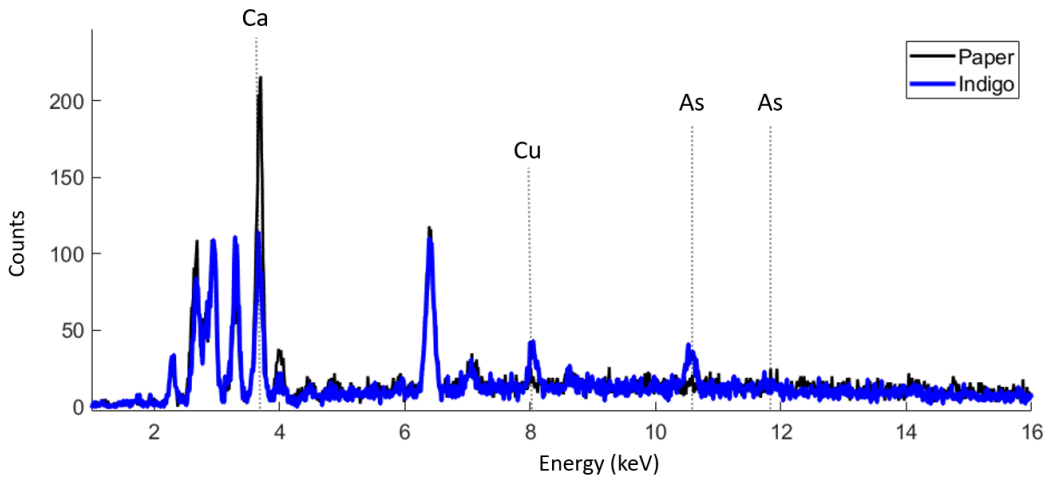

Fig.3.25 XRF spectrum for the blue regions in Or 16769, suggesting only impurities of copper and arsenic exist within an indigo mixture.

3.2.2.2 Prussian Blue

Prussian Blue is a dark blue synthetic pigment which was first mentioned in the early 18th century (Bartoll, 2008) and subsequently used in European paintings as a replacement for other more costly blue pigments such as natural ultramarine. It is produced as a result of the oxidation of ferrocyanide which consists primarily of iron (Fe). Prussian Blue has a history of use in different parts of the world, but in Asia, it was adopted and used most by artists in Japan during and after the Tokugawa (Edo) period (FitzHugh et al., 2003), which ended in the mid-19th century. It is documented that Prussian blue may have been largely imported to Japan by Dutch traders, and as plenty of the main trade routes to Japan from Europe required travelling through maritime Southeast Asia, it seems reasonable to assume that Prussian blue may have made its way into the hands of Southeast Asian artists in the 18th and 19th centuries. Furthermore, in the early to mid-19th century, Prussian blue was regularly manufactured from, and traded to China (Wang, 2023) from locations as far as London, making it almost expected to be encountered within maritime Southeast Asia.

Javanese Manuscripts

Evidence of Prussian blue exists on five different Javanese manuscripts where it can be identified through comparisons with simulated KM mixtures using known reference spectra. MSS Jav 28, MSS Jav 67, MSS Jav 68, Or 15932 and Or 15026, all possessed spectra indicating the use of Prussian blue, where well-fitting mixtures can be created to fit with VNIR mean cluster spectra from each manuscript using a Prussian blue reference with an added paper component for each manuscript (Fig 3.27), which only deviates towards the infrared most likely due to the high absorption of the Prussian blue references.

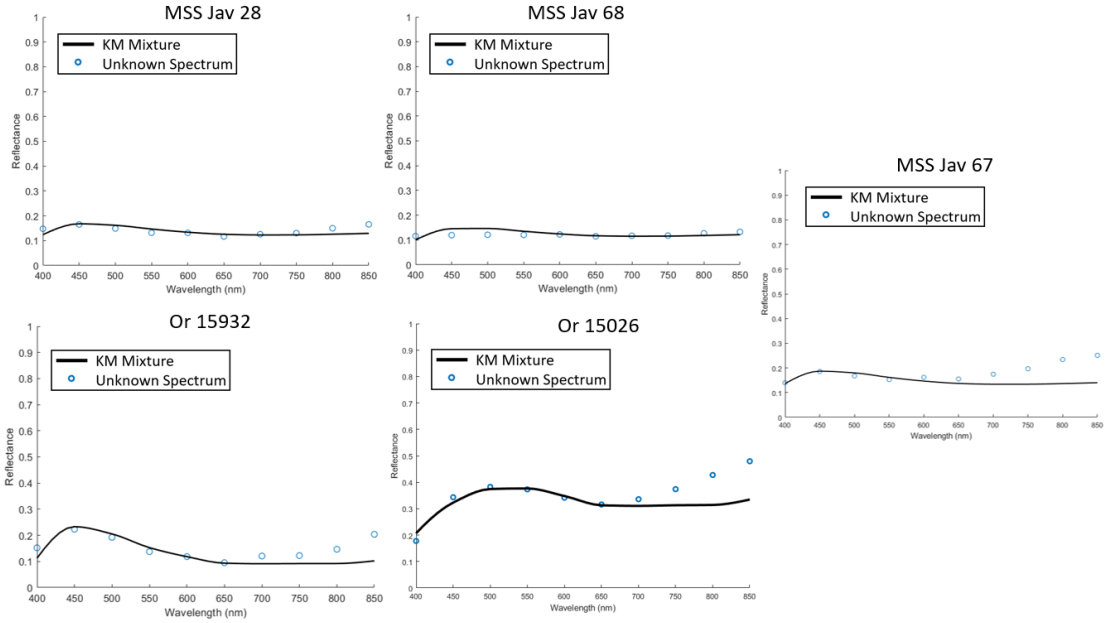

Fig.3.26 Multiple graphs showing KM mixtures for Prussian blue and paper in numerous Javanese manuscripts.

While this is a good indicator of Prussian blue being used over these many manuscripts, due to its high absorption and lack of clear spectral features, additional confirmation of the pigment was required through other means so that a conclusive identification could be performed. Additional evidence in favour of Prussian blue was collected via the use of XRF which can confirm the presence of elemental energy peaks higher than the paper for iron (Fe) at ≈6.4kev and ≈7kev, however, the pigments usage was fully confirmed using ER-FTIR (Fig 3.28). ER-FTIR spectra recorded for MSS Jav 68 and Or 15026, both provided IR spectra for the blue areas which possessed large absorption features at approximately 2085-2090cm^{-1}, corresponding reasonably well with the typical vibration of the carbon-nitrogen bond, which can often exist shifted from literature values to cover a region from approximately 2074-2100cm^{-1} (Petroviciu et al., 2019; Zuena et al., 2020); normally expected to be a consequence of spectral shifting due to the substrate, as was mentioned in Zuena et al., (2020).

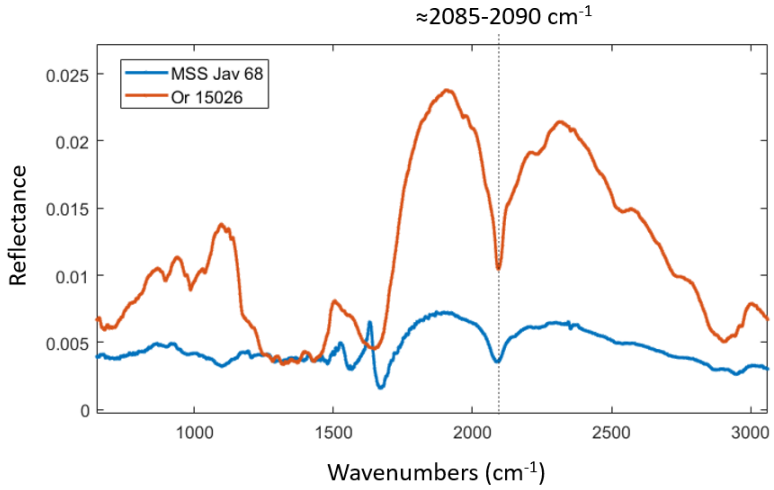

Fig.3.27 ER-FTIR spectra for MSS Jav 68 and Or 15026 which possess absorption features at ≈2085-2090cm^{-1} corresponding well with Prussian blue.

As the VNIR spectra for the other manuscripts suspected to use Prussian blue belong to the same or similar pigment mixture cluster groups as the dark blue recorded on both Or 15026 and MSS Jav 68, it was assumed that the blue used in all five manuscripts is likely Prussian blue. An interesting point to note for these manuscripts though is that the artistic styles and origins are quite similar. In the case of MSS Jav 28 and MSS Jav 68 for example, each dated to 1804 and 1805 respectively, the styles are almost identical, and as both manuscripts have the same source (Gallop, 2022), having been received by Col. Colin Mackenzie in 1812, and belonged to the 1822 Mackenzie collection, it comes as no surprise that these two manuscripts would contain the same types of pigments. Or 15932 and MSS Jav 67 are alike in style but are independent of the previous items MSS Jav 28, and MSS Jav 68, suggesting that even though the materials used are similar, their illustrations and illumination may have been performed by an alternative artist or artistic

school. It is interesting though that a similarly styled manuscript to both Or 15932 and MSS Jav 67 is Mackenzie Private 42, where the size and shape of the manuscript, and illustrations of human and animal characters are very similar, but for this manuscript, indigo is the only blue pigment confirmed. Or 15026 (dated to 1861) is the main standout manuscript here however, where the Prussian blue is mixed with considerable amounts of barium white to achieve a highly reflective VNIR spectrum in comparison to other manuscripts, clearly indicating a difference in implementation of colour for later Javanese manuscripts.

Malay Manuscripts

Two Malay manuscripts were found to possess Prussian blue after performing cluster grouping on the Javanese manuscripts, Or 13295 and Or 15227. Or 13295 uses Prussian blue both as a dark blue and a lighter shade achieved by mixing varying amounts of lead white, supported not only by the presence of iron in the XRF spectrum (seen in Fig 3.18) but also by a well-fitting VNIR spectrum created using KM modelling (Fig 3.29).

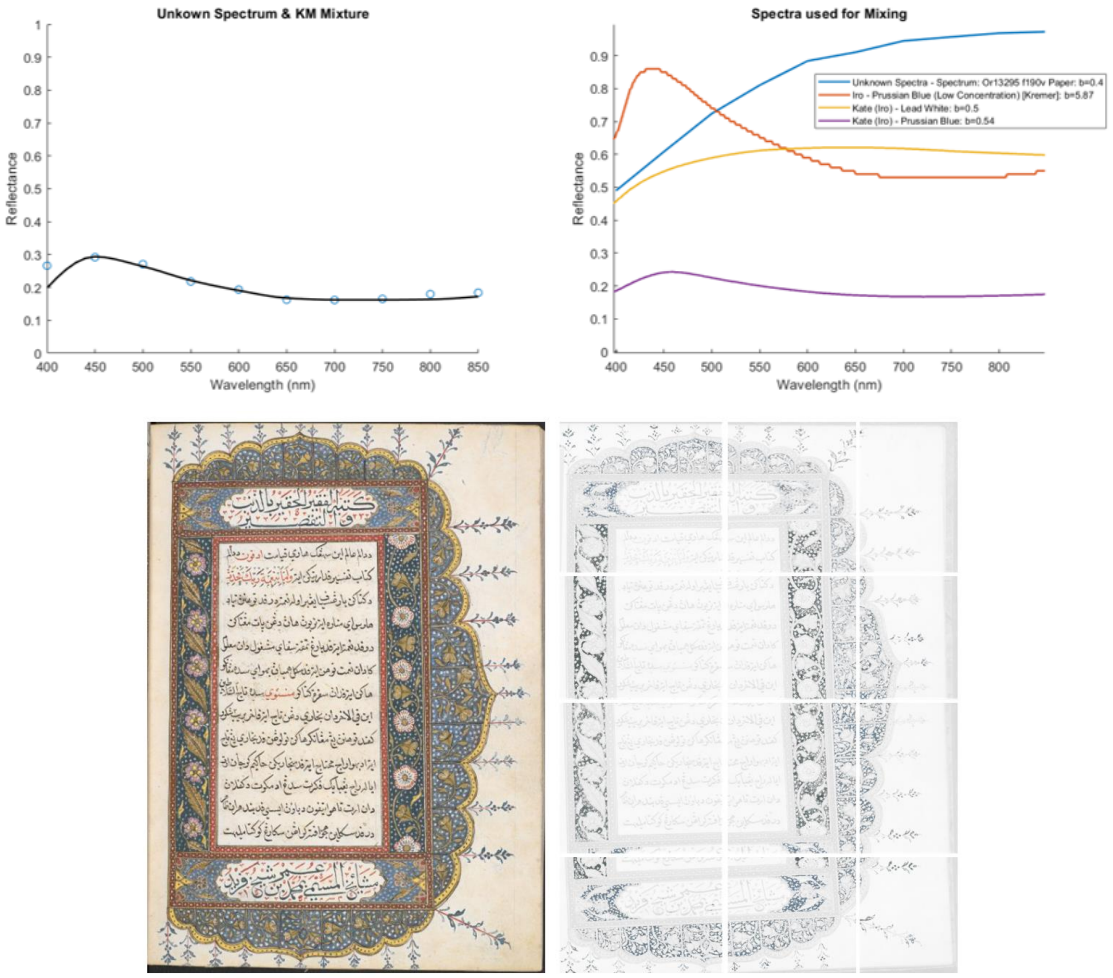

Fig.3.28 (Top) Well-fitting KM mixture of Prussian blue and lead white for the blue regions in Or 13295 f190v. (Bottom) Group cluster map for Prussian blue distribution on f190v after grouping Javanese clusters only. Note the preference fir the darker Prussian blue which was more common in Java.

As the graph shows, a good fit can be achieved for the Prussian blue when using paper, lead white, and two references for Prussian Blue, one with low concentration and another with high concentration, to help account for KM fitting difficulties which can occur when highly absorbing pigments are used. In contrast to Or 13295, another manuscript Or 15227, made much later in the second half of the 19th century, uses different thicknesses of Prussian blue directly onto paper to achieve different hues. This can be demonstrated well within the rubrication on folio f193v, shown in Fig 3.30, where the lighter blue central circle is more thinly applied than the stem-like areas around the outside, and therefore has a reflectance spectrum of higher intensity. This result is interesting as it suggests that even though the same materials are used, the approach to using them can vary depending on the time period or location.

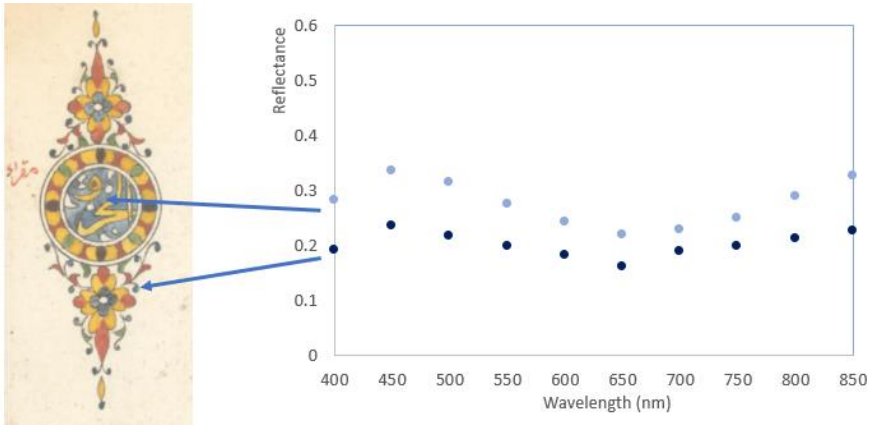

Fig.3.29 Image and graph showing the use of Prussian blue in different concentrations to produce darker or brighter blue shades in Or 15227 f193v

3.2.2.3 Ultramarine

Ultramarine is a blue pigment produced by grinding lapis lazuli and has seen widespread use throughout history. It has been used for hundreds of years, being used by artists in Europe during the Renaissance in the 15th century (Barnett et al., 2006; Dooley et al., 2014), though in Asia, ultramarine had been used in Chinese paintings and Indian murals much earlier (Plesters, 1966). Ultramarine was commonly transported from locations such as modern-day Afghanistan, with evidence of its use as early as the 6th-7th centuries, and while it is largely implied that natural ultramarine was likely always sourced from central Asia, there are other modern sources of lapis lazuli which exist today within southeast Asia (AGTA, 2023), potentially implying that by the 18th-19th centuries, more local sourcing of ultramarine may have been available. In the early 19th century, the synthesis of artificial ultramarine was developed (Hamerton et al., 2013), allowing for its mass production and causing a change of use from natural sources to artificial ones. If ultramarine is detected there is a chance for either artificial or natural variations to appear, as certain manuscripts can be dated to and beyond the early 19th century. With the data collection

performed throughout this thesis, ultramarine can be identified by VNIR reflectance where it has distinctive absorption features between 400-850nm. Furthermore, it can be detected using XRF, as it has the chemical formula $Na_{6-10}Al_6Si_6O_{24}S_{2-4}$, and the range of the Bruker CRONO can allow for weak detection of silicon (Si). Throughout the collection, ultramarine was detected in two late Javanese manuscripts, one Bugis manuscript, and the Lombok manuscript, all of which had dates ranging from the mid-18th century to the late 19th century.

Javanese Manuscripts – Or 15026 & Or 9333

In Or 15026 (1861), ultramarine is used for multiple Wayang characters. On folio f192r, a character can be seen illuminated with blue, for which PRISMS and XRF provided spectra indicative of ultramarine. For the VNIR, KM mixture modelling was used to produce a well-fitting ultramarine mixture to the mean cluster spectrum as shown in Fig 3.31. Furthermore, in addition to the VNIR, XRF spectra also showed the presence of both silicon (Si) and sulphur (S), providing further evidence of ultramarine (see lower figure in Fig 3.3). Ultramarine also appears in other parts of the manuscript but seems reserved only for the characters instead of any illuminated double frames, where Prussian blue is used instead. This perhaps suggests that the Wayang characters are treated differently due to their importance in Javanese culture, or that more than one artist illustrated the various areas, perhaps even at different times. In Or 9333 (late 19th-early 20th CE), VNIR reflectance spectra can be found in multiple blue regions which have features also closely matching ultramarine. For example, in folio f1r, a KM mixture using paper with ultramarine produces a well-fitting spectrum, as illustrated in Fig 3.32. There is however significant deviation from 700nm onwards, but this could be due to the change in the substrate between the reference material used in the KM mixture and the paper in Or 9333, where there is greater absorption in the blue/green parts of the spectrum, causing a closer fit between 400-550nm while overestimating infrared where ultramarine becomes more transparent. Interestingly, for the same reason, the clustering also completely segregates both Or 9333 and Or 15026 prior to grouping. The use of ultramarine in these Javanese manuscripts may seem sensible in terms of the availability of pigments during the dates of creation of these manuscripts being the mid-late 19th century, as during this time synthetic ultramarine may have been available. However, it is odd for these manuscripts that any ultramarine is found at all, considering that most items within the collection tend to use indigo or Prussian blue. It is important to note however that both of these manuscripts are very unique in their design and do not match well with many other artistic styles due to them both being created at much later dates (Or 15026: 1861, and Or 9333: late 19th-early 20th century). The presence of ultramarine in these manuscripts therefore clearly shows that it used by some artists from the mid-late 19th century onwards, which could suggest that it may only have been used in Java after the mass production of synthetic ultramarine had begun.

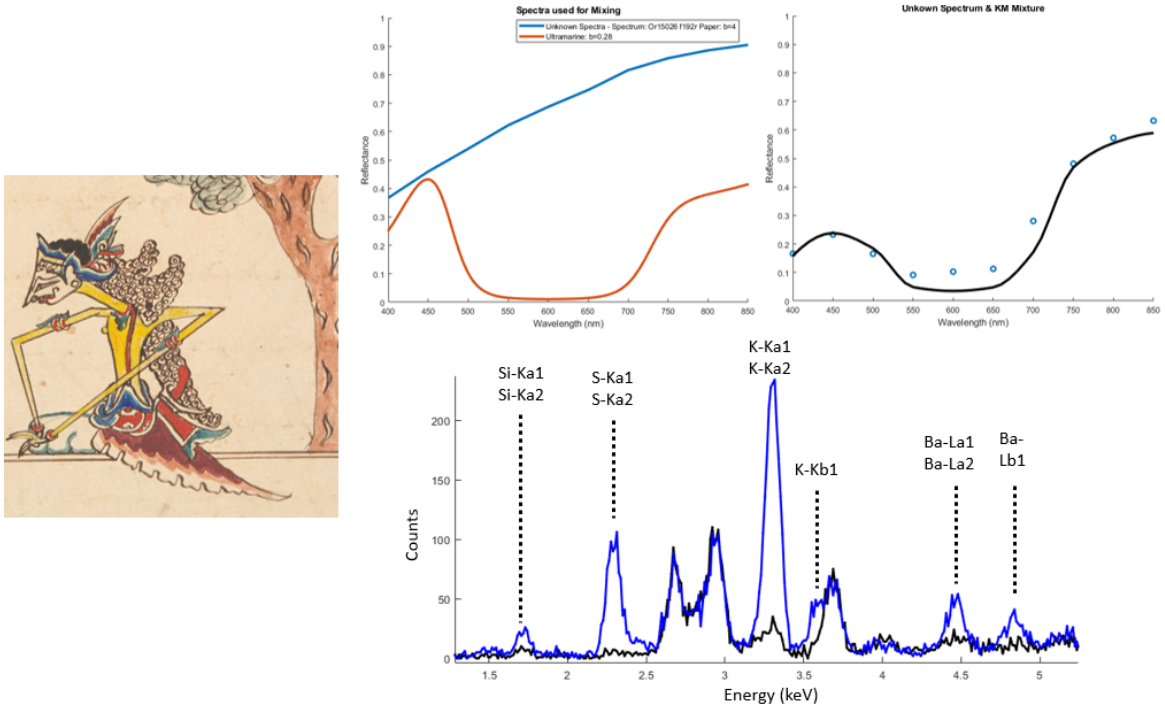

Fig.3.30 KM fitting graph and XRF spectrum showing the existence of ultramarine blue in manuscript Or 15026

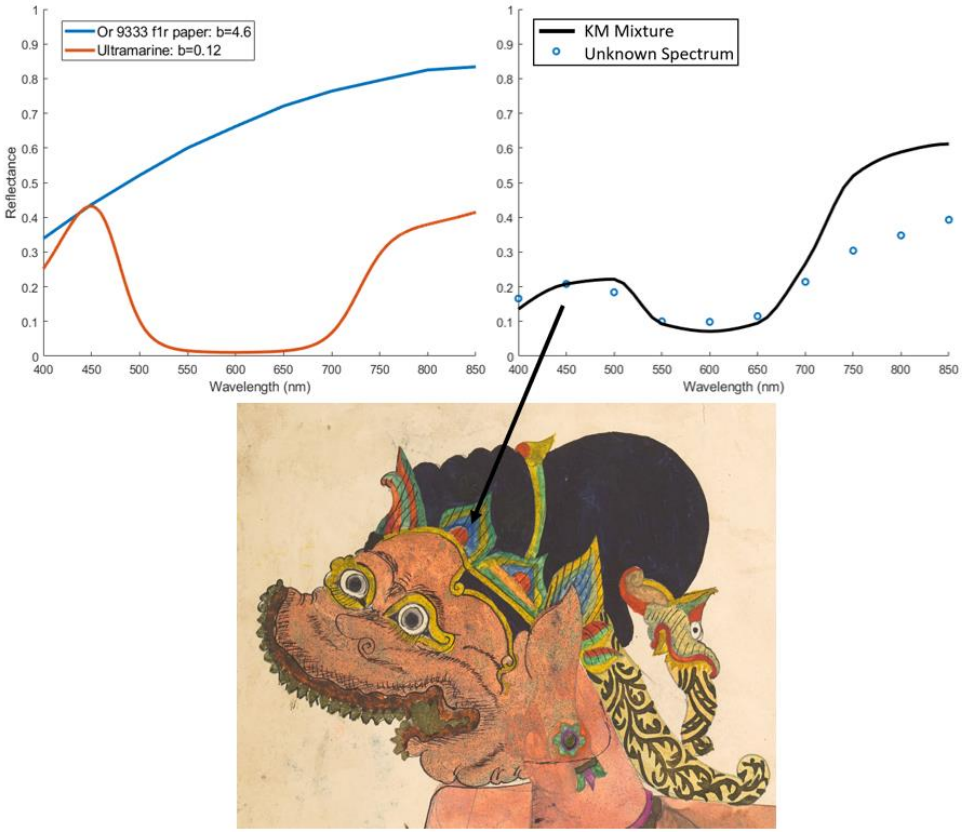

Fig.3.31 Graphs showing the KM mixture for ultramarine blue with paper, which is reasonably well-fitting for a blue region in Or 9333. There may however be other materials due to the slight differences in spectral shape at 500nm.

Lombok Manuscript – Or 15924

Or 15924, a scroll from Lombok, seemingly uses ultramarine in the bordered parts of the paper. When extracted PRISMS spectra taken from the blue regions are also analysed using KM mixture modelling, the best fitting mixtures are always found when using ultramarine as a component, either with the paper or also with the addition of lac, though this could not be confirmed. This can be seen in Fig 3.33, where two different mixtures are demonstrated showing well-fitting reflectance spectra. As the scroll is dated to the late 19th-early 20th centuries (Ricklefs et al., 2014), it is reasonable to assume that by this time, synthetic ultramarine blue may have been common throughout maritime Southeast Asia, and therefore by extension in Lombok. It also connects well with the other ultramarine usage being linked to later Javanese manuscripts, suggesting that over the entirety of the archipelago, ultramarine may have become more common towards the early 20th century.

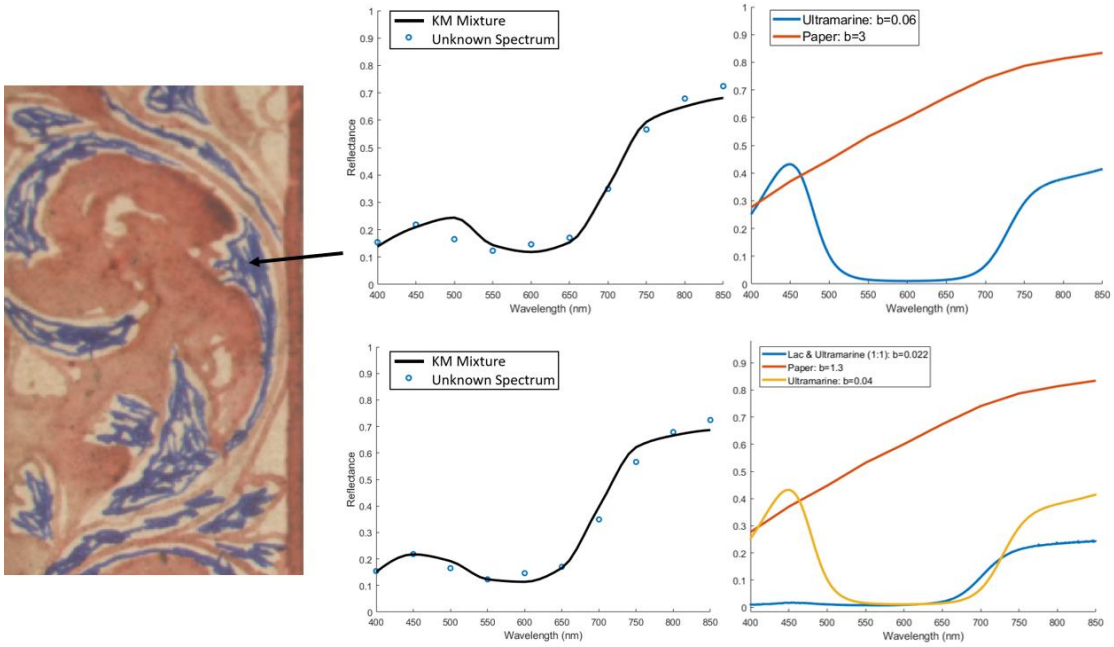

Fig.3.32 KM fits showing well-fitting ultramarine mixtures for the blue in Or 15925. Note however that a mixture of ultramarine and lac produces an even more accurate fit, suggesting the use of this mixture.

3.2.2.4 Copper-based Blue (Potentially Azurite)

Evidence of a copper-based blue, potentially Azurite, exists within Or 9333, the latest manuscript in the collection. While azurite has been used in other parts of Asia, its existence in maritime Southeast Asian manuscripts during the 18th-19th centuries is rare. The potential presence of azurite could be estimated by extracting VNIR reflectance spectra and performing KM mixture modelling for the blue area, which indicated a well-fitting spectrum on folio f.5v when using paper and a sensible contribution of malachite. XRF spectra showed a strong presence of copper

(Cu), as demonstrated in Fig 3.34, but also silicon (Si), potassium (K), iron (Fe), mercury (Hg) and arsenic (As), suggesting other materials and impurities, assuming the signal is not sourced from elsewhere in the manuscript. While the presence of Hg and As can likely be attributed to the nearby yellow/orange and vermillion red overleaf, the presence of other elements could imply other blues, such as the iron-based (Fe) Prussian blue, or ultramarine which contains silicon (Si). However, as none of these other pigments match well with the VNIR spectrum, it is more likely that the presence of these elements is due to impurities within the pigment, as many of these elements can be attributed to other natural minerals sometimes found in azurite (Aru et al., 2014).

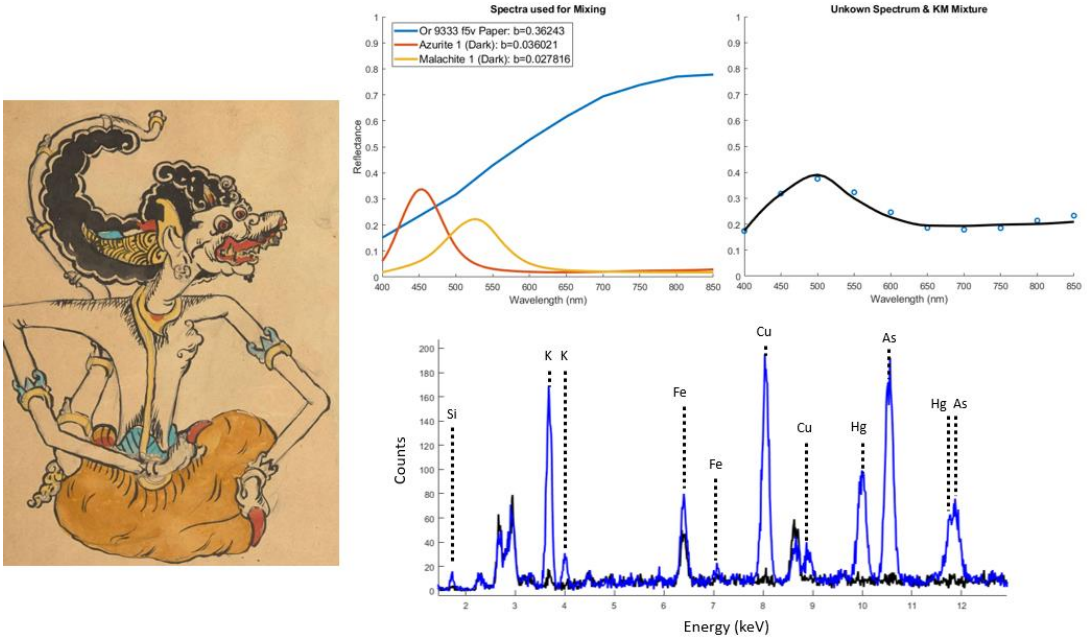

Fig.3.33 KM mixture showing that azurite and malachite together can create a good fit to the unknown spectrum and can be confirmed by the strong presence of copper in the XRF.

While the detection of copper in this manuscript does not confirm the presence of azurite, there is clearly a contribution of copper in the blue region. If this pigment is in fact azurite, this would come as a surprise considering its rarity as use as a pigment within maritime Southeast Asia. However, as this manuscript is quite different from the rest of the collection, it makes sense that there may be some differences in the implementation of different artistic materials. While the use of azurite in maritime Southeast Asia is not common, its presence could suggest a connection between both China or the Islamic world, as both cultures are documented to have used azurite in different artistic practises and manuscript illumination (Knipe et al., 2018; Ngan, 2018), although by the 19th century, its use would have been rare, as most regions would have likely used cheaper and more accessible alternatives in the same way that Europe had done so much earlier (Price, 2000). It is unlikely that azurite would have been sourced locally, as there is little evidence to suggest that azurite was a commonly occurring pigment within the archipelago, with only a select few mineral studies suggesting the contrary, e.g. (Willbourn, 1925), and even then it is still noted to remain uncommon.

3.2.2.5 Blue Mixtures

In one manuscript there appears a mixture of multiple pigments used to create a dark blue colour. While most blue pigments can be identified easily using VNIR spectra, the high absorption of this dark blue made the detection of materials very unclear. The dark blue, used in the hooves of Buraq Or 8154 f3v, appears to be very unique and unlike any other blue pigment or mixture encountered throughout the rest of the collection. When monitoring the group mean cluster VNIR reflectance spectrum of the dark blue hooves, it is not immediately clear which blue pigment the material consists of. Prussian blue appears to be the closest matching reference to the unknown mixture and can be confirmed in the blue tail after grouping the Javanese and Malay Prussian blue clusters, however, there is an increase in reflectance after ≈650nm for the hooves which could imply the presence of indigo or ultramarine. Attempting to use KM to fit for different mixtures yielded no useful results, as any number of different combinations of indigo, ultramarine, and Prussian blue in different concentrations all result in a fit which is close to the unknown spectrum, though a mixture of all three appears to fit best, as shown in Fig 3.35.

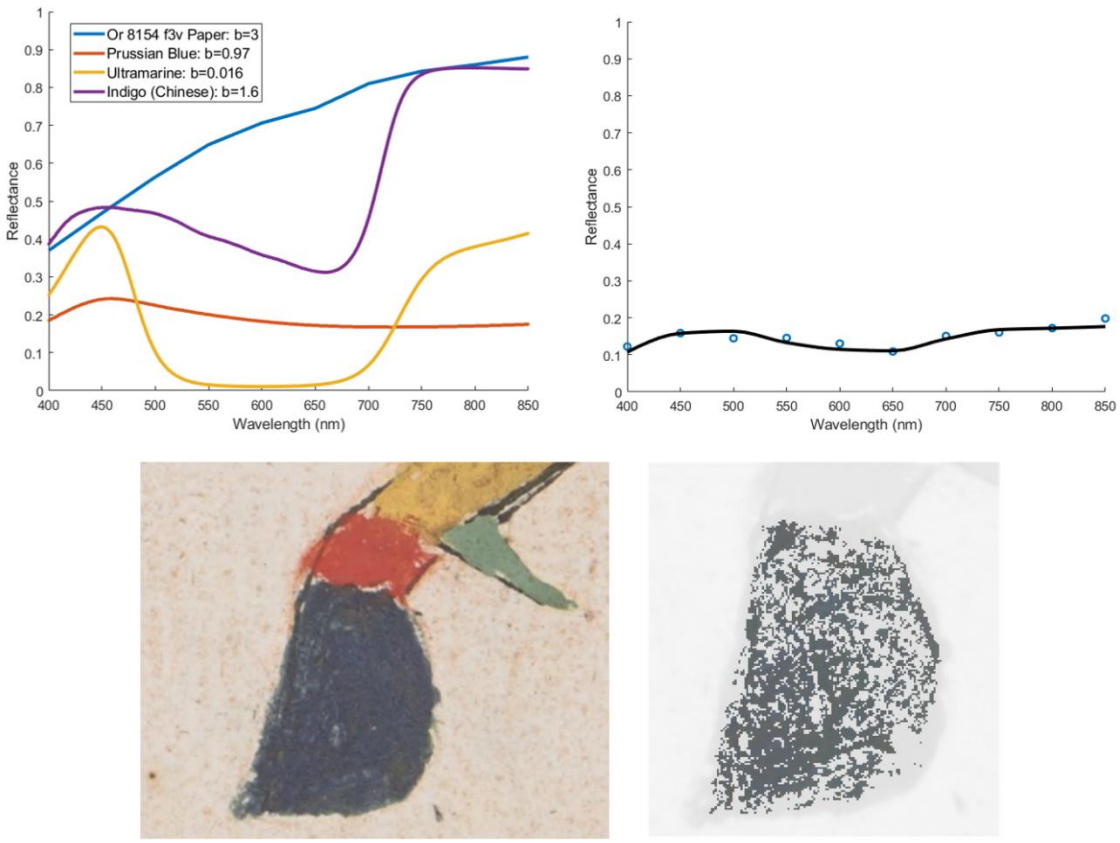

Fig.3.34 KM Mixture of indigo, Prussian blue, and ultramarine which creates a reasonably well-fitting spectrum to the unknown blue mixture.

As different shades of blue and a much darker, almost black, pigment can be seen within the hoof, it could be possible that there is a layering of one blue atop another which may have darkened over time. To further understand which constituent materials may be within the hoof XRF data was also recorded, however, even this did not provide any definitive evidence of exactly which

blue pigments have been used, as both consist of many elemental peaks which could be indicative of multiple materials (see Fig 3.36).

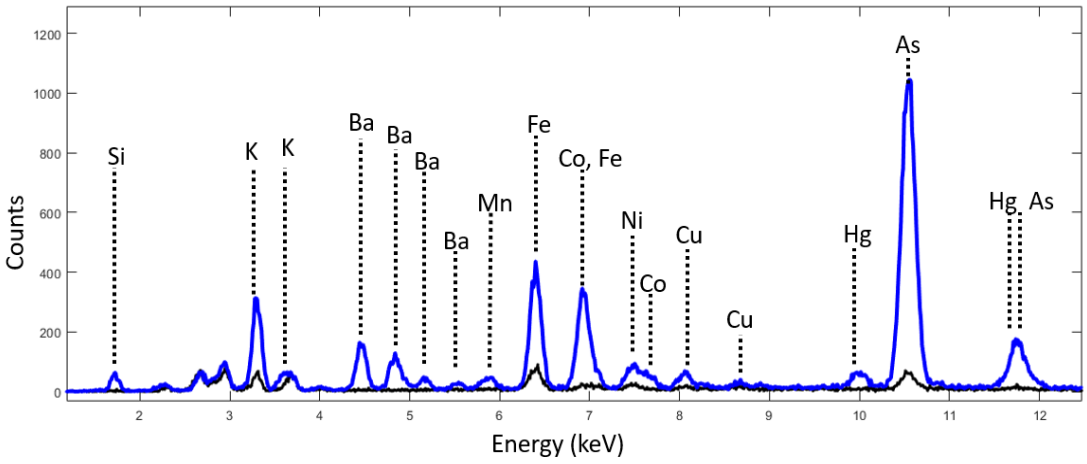

Fig.3.35 XRF spectrum for the dark blue hoof in Or 8154, which shows the presence of many different elements which could be attributed to different blue pigments.

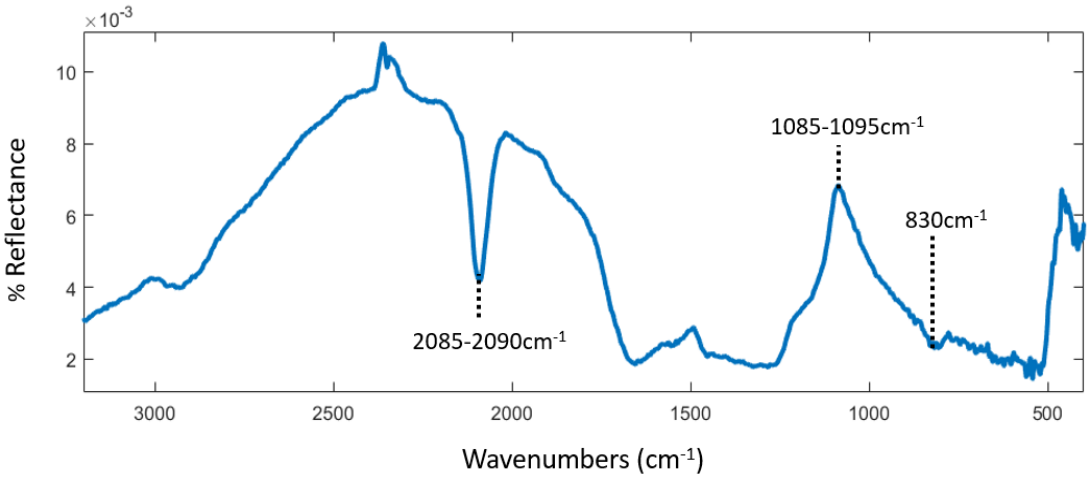

Fig.3.36 ER-FTIR spectrum indicating the presence of both Prussian blue and ultramarine in the dark blue hoof on Or 8154 f3v.

If anything, the XRF adds greater complexity, as the existence of many other elements only increases the number of potential components within the mixture or layers. As both XRF and PRISMS spectra are unable to provide any definitive results, ER-FTIR data for the navy hoof was analysed to investigate potential materials, the spectrum of which can be seen in Fig 3.37. The ER-FTIR spectrum possesses a peak at ≈2080-2090cm^{-1} which matches closely with that of Prussian blue seen earlier, which comes as no surprise considering it was the most likely pigment to be used due to its high absorption. There is also a broad feature at ≈1085-1095cm^{-1} which corresponds with the stretch of the Si-O-Si bond often found in ultramarine reference spectra and weaker bands at ≈830cm^{-1} which may further fit with the Si-OH stretch (Jensen, 2023) . It would seem appropriate then to assume that this dark blue mixture is a combination of at least both Prussian blue and ultramarine.

3.2.3 Red-Pink Pigments & Mixtures

Throughout the collection, numerous pigments and dyes are used to achieve different red colours, these include red ochre, insect/plant dyes, red lead, and most commonly, vermillion.

3.2.3.1 Vermillion

Vermillion, or Mercury Sulphide (HgS), is an inorganic red pigment found over much of the world. In Asia, vermillion saw widespread use in China for many artworks and other objects such as pottery and lacquerware (Kopplin, 2002), so much so that some shades of the pigment are referred to as Chinese Red. During the 18^{th}-19^{th} centuries, vermillion was used as the primary red pigment throughout most of the world, but was known to be a commodity often traded between China and Southeast Asia as early as the 15^{th} century via Malacca, one of the most important trade ports of maritime Southeast Asia located on the Malay Peninsula (Orillaneda, 2016). Due to this it would be an appropriate assumption to suggest that vermillion may have seen common use in maritime Southeast Asia. To confirm whether vermillion is used within the collection, XRF was primarily used alongside PRISMS reflectance spectra to both identify the chemical elements present within the red areas and perform a comparison of extracted VNIR spectra with references.

Javanese Manuscripts

Under closer inspection of the red-painted areas for each Javanese manuscript, most of the associated VNIR reflectance spectra appear to have only a single sharp feature that is indicative typically of pigments such as vermillion. Using a combination of PRISMS VNIR spectra and XRF, vermillion could be identified on every Javanese manuscript which implemented red into its palette, both independently of other mixture components and as a component in the creation of other colours such as pink or brown. Some example KM mixtures representing measurements taken from multiple red regions in different manuscripts can be seen in Fig 3.38. The multiple plots in Fig 3.38 clearly illustrate that using only the paper of the manuscript and vermillion as the two main components creates a virtual mixture which fits well in red areas over many manuscripts. The only deviations between the mixtures and the true extracted spectra appear towards the infrared parts of the spectrum, which can most likely be attributed to the spectrum following the reflectance intensities of different the paper substrate when the paint application is thin. With these results it appears clear that vermillion is most likely the pigment used in most manuscripts, however, this can be proven more conclusively when using XRF. Normalised and vertically displaced XRF spectra extracted from the same Javanese manuscripts show the presence of mercury (Hg) and Sulphur (S), as seen in Fig 3.39, providing further evidence of the red pigment in each manuscript being vermillion.

While for most Javanese manuscripts a simple mixture of vermillion with paper fits well to the extracted spectra, there are some circumstances in which the inflection point of a vermillion KM mixture does not completely fit with the extracted PRISMS VNIR spectra. This is not a surprise as KM mixtures created using pure references do not always fit perfectly with in-situ recorded reflectance spectra, however, there is also a chance that this spectral difference is attributed to the presence of red dyes or red lead in addition to vermillion (as discussed in more detail in 3.4.3.3 – 3.4.3.4).

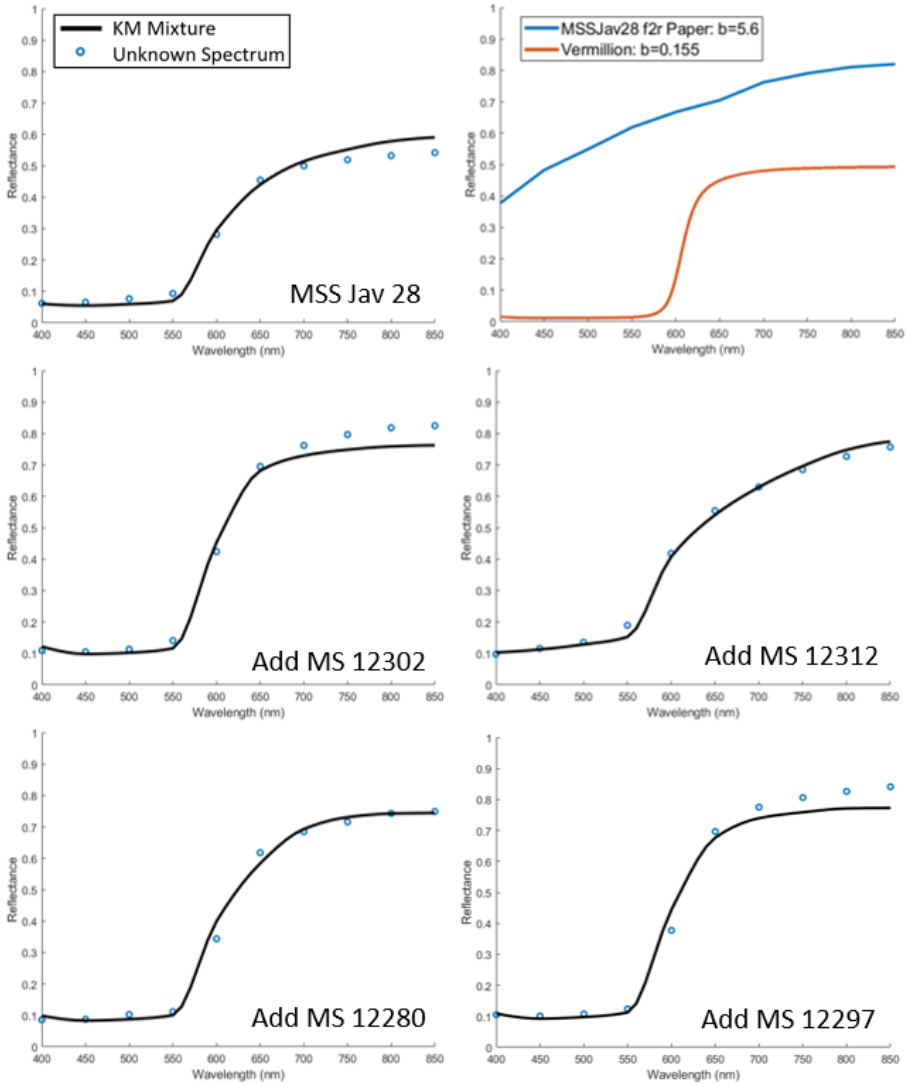

Fig.3.37 Multiple graphs showing KM mixtures for Vermillion and paper in 5 different Javanese manuscripts.

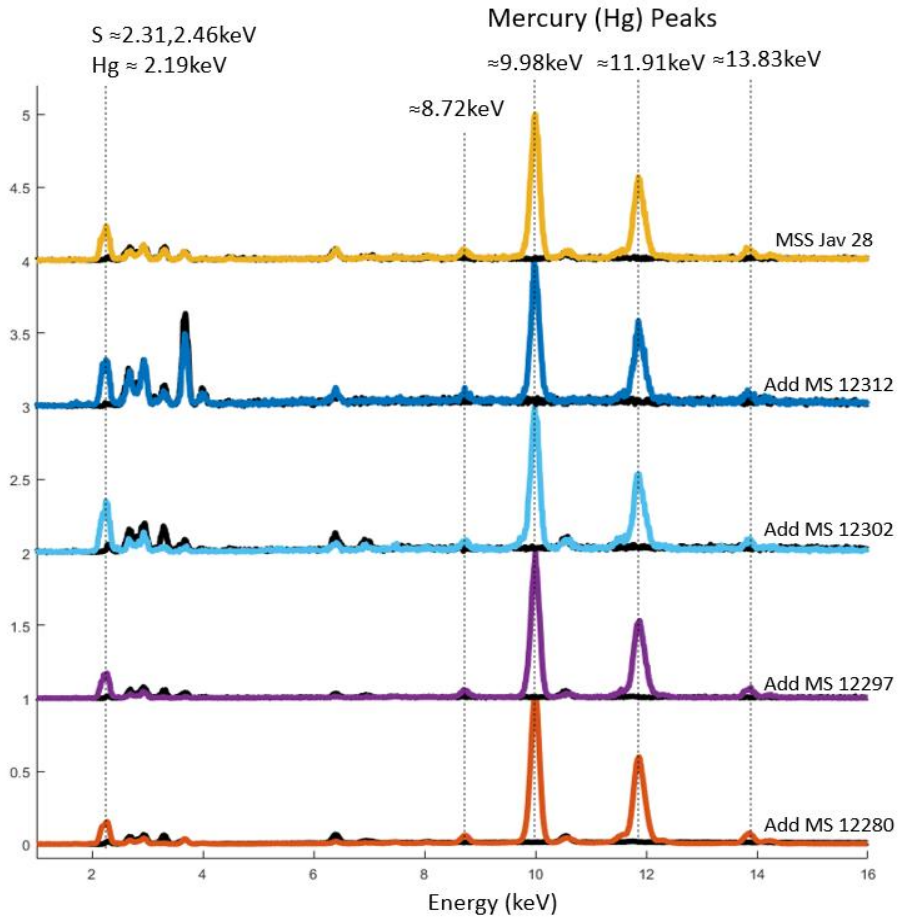

Fig.3.38 Multiple graphs showing XRF spectra for red areas (coloured) with paper (black) for 5 different Javanese manuscripts. The presence of Hg and S indicated the use of vermillion in all.

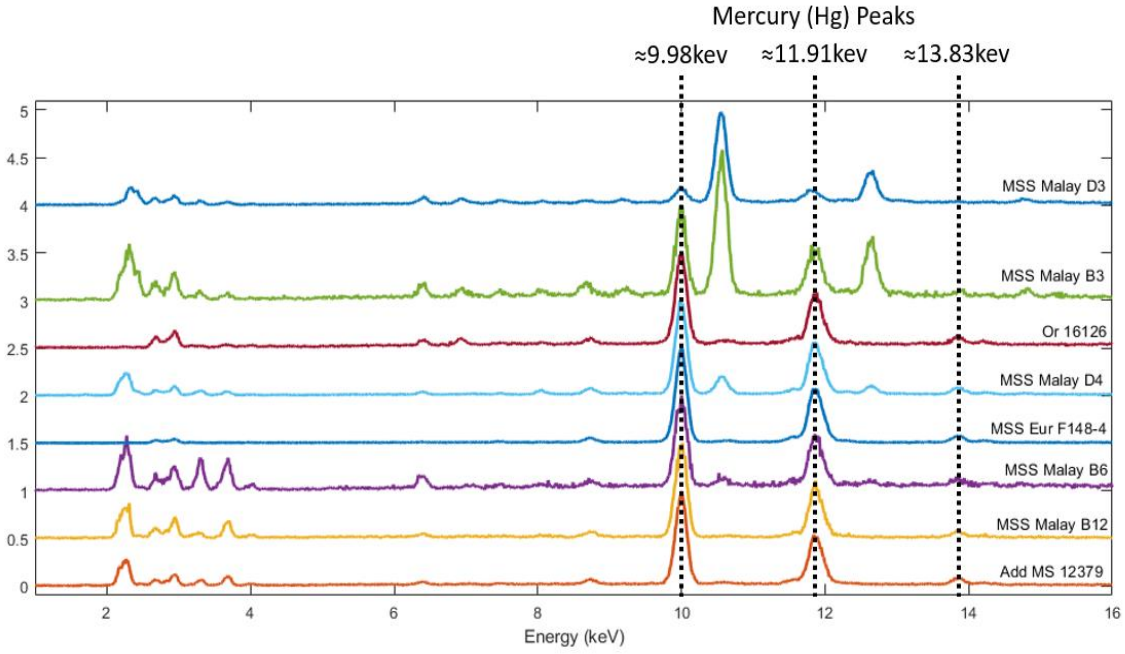

Fig.3.39 XRF spectra showing the presence of Hg peaks indicative of vermillion in multiple Malay manuscripts.

Malay Manuscripts

Within the collection every Malay manuscript uses vermillion as a red pigment in some form, however, its use is not exclusive of other red pigments, as can be demonstrated in some later sections where red ochre and dyes can both be found within different manuscripts. The presence of vermillion can be confirmed by the presence of mercury (Hg) and sulphur (S) in every manuscript where red was used, as seen in Fig 3.40, where different XRF spectra taken of red areas can be seen after normalising each spectrum and displacing them vertically for clarity. In all the XRF spectra, the typical mercury peaks can be seen, normally with additional sulphur peaks which vary in intensity depending on the manuscript, with the weakest S peak found in Or 16126 (1857) where only a trace is detected. In MSS Eur F148-4 (≈1811) however, no sulphur was detected by XRF at all even though the mercury was present as it has the three main peaks as illustrated in Fig 3.40. This is likely a result of performing XRF through a plastic film with which the letter was protected, as plastics can cause further absorption of X-rays at lower energies, therefore for this scenario, the pigment used is most likely still vermillion. In addition to vermillion, the red areas in some of the manuscripts appear to be mixed with other materials. The presence of both lead and iron peaks in multiple spectra may exist as a result of the mixing between vermillion and red lead, or red ochre respectively, a not uncommon approach to use if vermillion became too expensive or if an artist worried about its toxicity, though this is unlikely during this time period. It could also be the case that the vermillion is mixed with lead white to achieve different hues. The detection of vermillion within all the Malay manuscripts should come as no surprise considering the known trade activity of the pigment in and out of southeast Asia through ports on the Malay peninsula. It is interesting however that even though vermillion should have been quite accessible to artists within the Malay peninsula, it still appears to be regularly mixed and used with other red pigments such as red ochre or perhaps red lead. This suggests that in certain regions or times, certain artists may have added other red pigments to vermillion intentionally.

Bugis & Acehnese manuscripts

With the grouping of clusters performed for both Malay manuscripts and Javanese manuscripts, it was found that throughout the rest of the collection, most red pigments had been vermillion, a similar result to what was shown previously in chapter 2. With these results being confirmed by XRF, the final pigment mixture table was updated with the presence of vermillion found in almost every manuscript, however in addition to vermillion there were also other red pigments and pink colours not placed into the vermillion clusters. These other red hues and their pigments are covered in the following subsections.

3.2.3.2 Red Ochre

As briefly mentioned in the previous section, there is evidence of red ochre to have been used within the collection. Red ochre is an inorganic red earth pigment which is produced as a result of the oxidation of iron (Fe) into iron oxide (Fe_2O_3). It has been used for thousands of years and is well documented to have been used for cave paintings which still exist today due to its chemical stability. Red ochre has seen consistent use in maritime Southeast Asia for many different purposes, with evidence of its use being traced as far back as 42,000 years ago in Timor-Leste, where archaeological excavations have uncovered marine shell beads which possessed trace amounts of red ochre (Langley & O'Connor, 2018). Therefore, it seems reasonable to expect that red ochre may appear in use within the British Library's maritime Southeast Asian collection. As its only inorganic constituent chemical element is iron (Fe) it can be identified using XRF quite easily, as such this was the primary source of evidence for identifying it in this collection.

As the preference for red pigment in Southeast Asia appears to be vermillion, red ochre was found only a few times within the collection, with no clear examples in the Javanese manuscripts at all. Furthermore, for most manuscripts in which it was determined to exist, red ochre was only ever mixed with vermillion, as XRF peaks for iron in red regions normally accompanied additional peaks for mercury and sulphur. In some regions, extracted PRISMS spectra which were expected to contain an influence from both vermillion and red ochre matched best with KM mixtures of only vermillion, which could suggest that detected iron content, which could be seen in XRF, was a result of traces of ochre or other materials being mixed into the pigment, or that layers of ochre may have been laid before applying vermillion, a question which is explored more in this section for MSS Malay B6.

Malay Manuscripts – MSS Malay B6

Within the Malay collection, there is evidence for red ochre to exist in only one manuscript as a pure pigment. In MSS Malay B6, the VNIR reflectance and XRF spectra extracted from a red region in folio f1v can be seen to fit well with red ochre and contain strong peaks for Iron (Fe) respectively, as shown in Fig 3.41. In addition, peaks for lead and trace amounts of mercury can be found, suggesting that red lead, lead white, or some vermillion may be mixed into the red ochre, however, it is difficult to confirm which.

For MSS Malay B6, it is interesting that red ochre and vermillion are used interchangeably as the main red pigment depending on which page is illuminated. For example, in f1v red ochre is used as the main pigment, whereas on the following page, f2r, vermillion appears to be used the most. It is unclear why an artist would choose not to use vermillion on both opening pages however one reason for this if the choice was not intentional could simply be that the manuscript was unfinished during its first illumination and either the vermillion or red ochre were added later.

It could also be possible that red ochre was used first to denote where scarcer and more brilliantly coloured vermillion paints should be applied later. This is somewhat supported by the fact that fitting KM mixtures to the PRISMS spectra of the vermillion in f2r shows the most accurate fits when including red ochre as a component, as shown in Fig 3.42. Furthermore, XRF spectra recorded for the red regions on folio f2r also confirm the presence of iron in lower concentrations than mercury (also in Fig 3.42), indicating a higher vermillion content or that the vermillion is layered above the red ochre, which seems more reasonable considering that the colour and VNIR spectrum both indicate an increased presence of vermillion.

Sulawesi (Bugis) Manuscripts

A single Bugis manuscript, Add MS 12372, was found to have evidence of possessing red ochre mixed with vermillion in the red regions on f68v. This could be seen in the VNIR spectra where a KM mixture of both vermillion and red ochre was able to provide a virtual spectrum which matched closely with the PRISMS spectra extracted from one of the illuminated red circles seen on folio f68v red region, as shown in Fig 3.43. Furthermore, the XRF spectrum for the same red region was shown to possess both mercury (Hg) and sulphur (S) peaks, as expected for vermillion, but also had peaks for iron (Fe) which were higher in counts than the paper, further suggesting the existence of red ochre.

Sumatran (Acehnese) Manuscripts

Red ochre could also be found in Acehnese manuscripts, where Or 15979, Or 16034, and Or 16035 all possess evidence of red ochre being mixed with vermillion. In both Or 16034, and Or 16035, vermillion is the main pigment used, however, XRF spectra show a presence of iron, suggesting that some red ochre may be mixed into the vermillion, as shown in Fig 3.44. In Or 15979, there appear to be regions where red ochre is found in much greater quantities, and even in some areas has taken an appearance closer to that of brown ochre. The red ochre in these regions can also be detected due to the higher iron content found in XRF spectra, however, it can be confirmed more clearly through the use of VNIR spectra, as shown by creating a KM mixture to fit a PRISMS spectrum for a light reddish-brown area on Or 15979 f2v (Fig 3.45).

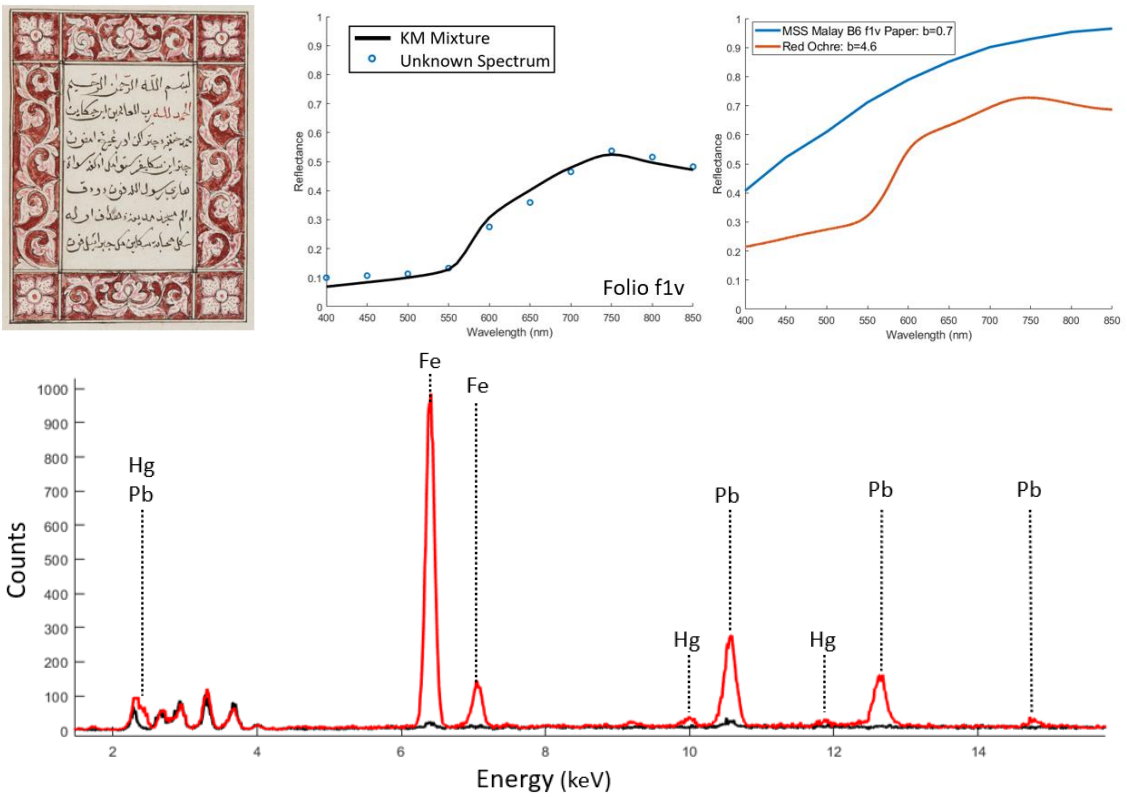

Fig.3.40 Graphs showing the presence of red ochre in MSS Malay B6 f1v. Top: RGB image and KM mixtures showing a well-fitting mixture of red ochre and paper. Bottom: XRF spectrum indicating a strong presence of iron (Fe) with other elements.

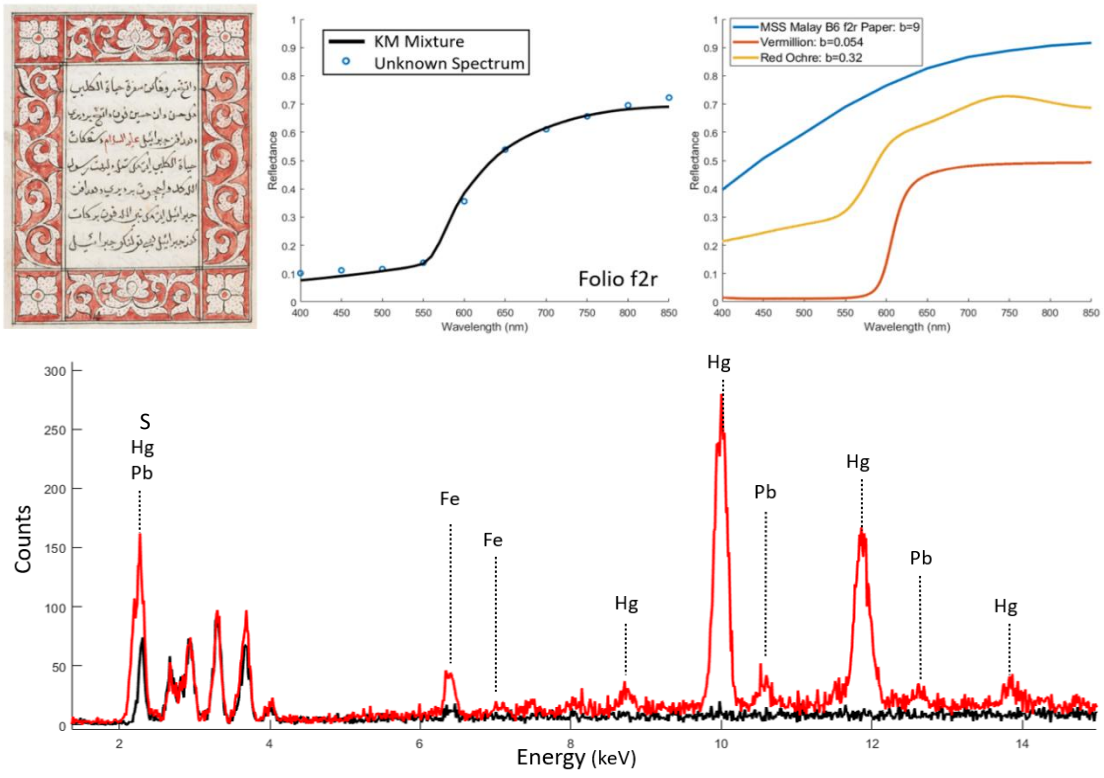

Fig.3.41 Graphs showing the presence of vermillion and red ochre in MSS Malay B6 f2r. Top: RGB image and KM mixtures showing a well-fitting mixture of red ochre, vermillion, and paper. Bottom: XRF spectrum indicating a strong presence of mercury (Hg) and sulphur (S), along with iron (Fe) and other elements.

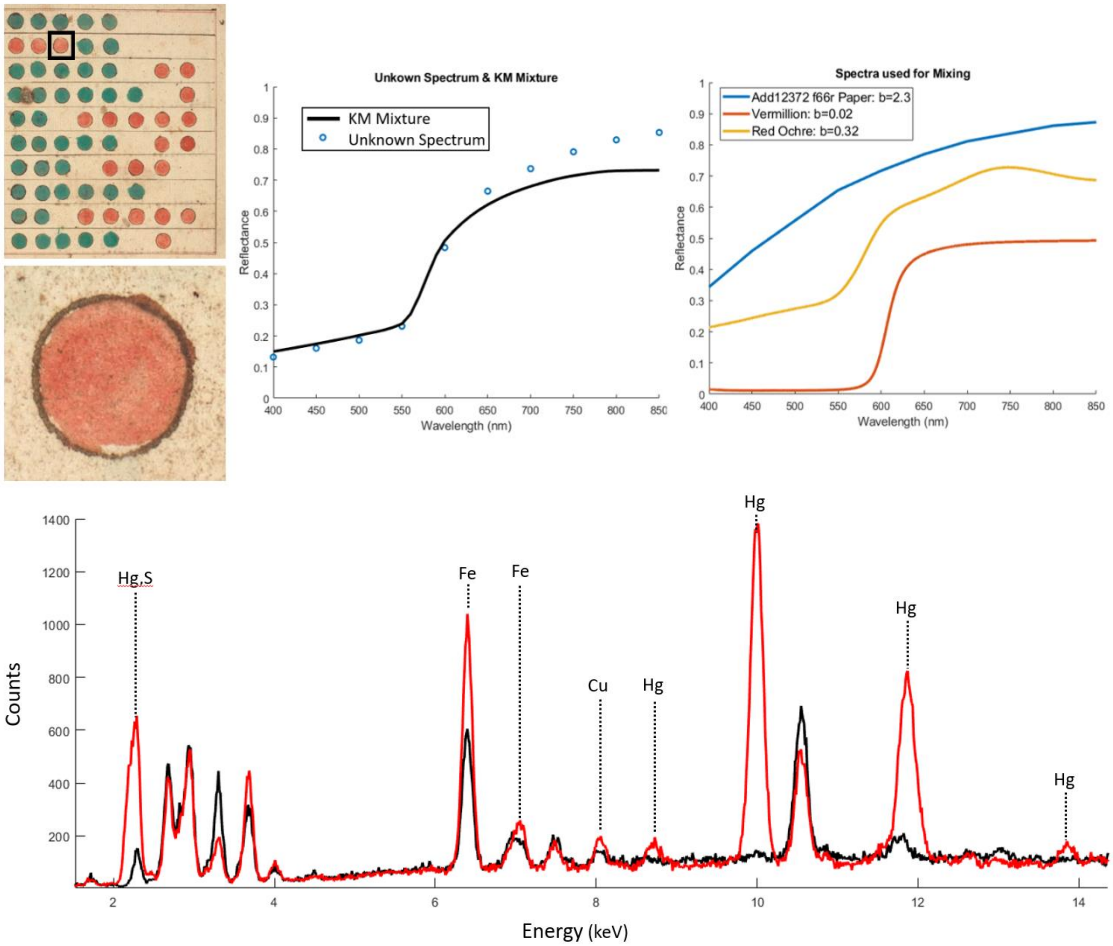

Fig.3.42 Graphs showing the presence of vermillion and red ochre in Add MS 12372 f68v. Top: RGB image and KM mixtures showing a well-fitting mixture of red ochre, vermillion, and paper. Bottom: XRF spectrum indicating a strong presence of mercury (Hg), sulphur (S) and iron (Fe).

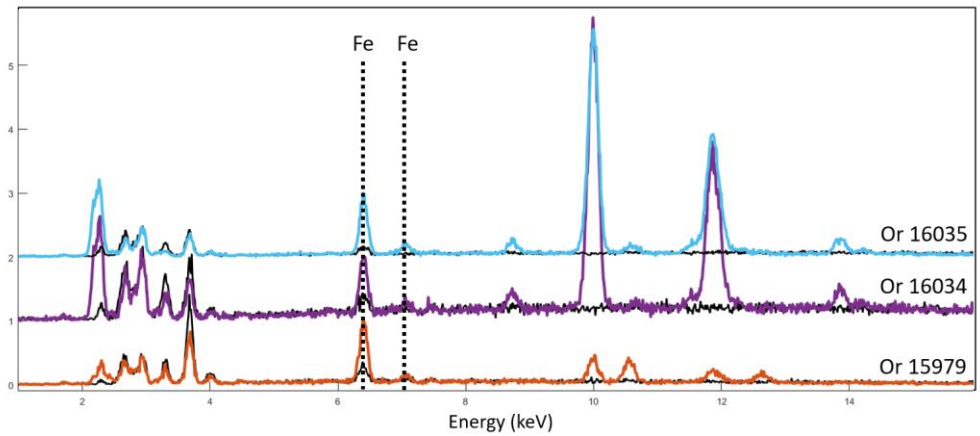

Fig.3.43 XRF spectra showing the presence of iron (Fe) in three Acehnese manuscripts, indicating the use of red ochre.

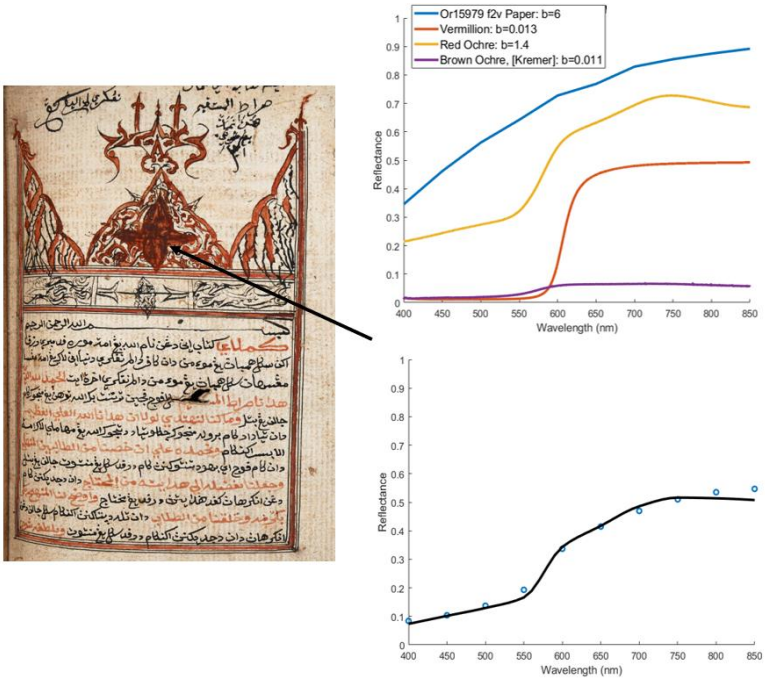

Fig.3.44 KM mixture showing a well-fitting spectrum using paper, vermillion, red ochre, and brown ochre as components in manuscript Or 15979, folio f2v.

3.2.3.3 Red Dyes

Throughout the collection, there are numerous manuscripts which implement red dyes either in addition to or exclusive of other red pigments. For the red dyes, clustering of PRISMS data was very important to be able to gather enough pixels to show clearly and confidently the presence of indicative absorption features. Therefore, for many of the spectra extracted for this section, the cluster maps are included to illustrate the group of pixels used to generate the mean spectrum for KM mixture modelling.

Javanese Manuscripts – Or 15932

Beginning with the Javanese manuscripts, aside from some red dyes being potentially mixed into different red colours (as is discussed later in section 3.4.4.6), there is very little evidence to suggest that red dyes are used in many Javanese manuscripts, with most red colours tending to use mainly vermillion. Two notable exceptions to this however are Or 9333, where red dyes are used in numerous Wayang characters, and Or 15932 where red dyes are implemented in abundance over many pages. Or 15932 notably uses dyes in two main colours over the manuscript. The first, which will be covered in this section, is a typical red dye colour which takes the form of burgundy or magenta on different pages. The second, which is covered in more detail in section 3.3.4.8, is a mixture where red dye appears to be mixed with a blue, suspected to be Prussian blue, to create a purple colour on folio f72v. Extracting VNIR PRISMS mean cluster spectra from the dark reddish robe of the character on Or 15932 f17v provided a spectrum which possessed absorption features indicative of red dyes from approximately 400-550nm. For this reason, KM

mixture modelling was used to investigate whether a mixture of a reference red dye with paper could create an accurate fitting virtual spectrum to fit the unknown red. The KM mixture created can be seen in Fig 3.46, with a pigment mixture map for the red dye, which shows that the best-fitting red dye was found to be cochineal carmine, though with deviation from the unknown mixture after approximately 650nm onwards. To be sure of dye being used, FORS was also used to provide higher spectral resolution, of which the resulting spectrum and new KM fit can be seen in Fig 3.46 and helps to confirm the presence of red dye. It is important to however that even though cochineal carmine was used for KM modelling, this does not confirm its use. However, the spectrum did possess absorption features which matched closely with a reference for cochineal carmine (containing absorption features at 525nm and 568nm), suggesting the use of insect-based dye. Or 9333 uses multiple dyes applied straight onto paper, where evidence can be found in favour of them existing to create a purple colour. The pink mixture found on f10v is discussed further in section 3.3.4.7, as it seems to be applied exclusively mixed with a white pigment. The purple also applied on f10v is also a mixture which is suspected to be mixed with blue, and so can be seen later also.

Fig.3.45 RGB image and true colour cluster map showing the use of red dye in Or 15932 f17v with multiple graphs showing KM mixtures for PRISMS (top right) and FORS (bottom) indicating the use of cochineal carmine.

Malay Manuscripts – MSS Malay B6 & MSS Malay B12

Two Malay manuscripts have evidence of red dyes being applied seemingly straight onto paper. These can be found in MSS Malay B3, and MSS Malay B6, where indicative absorption features can be found at approximately 400-550nm in the extracted PRISMS spectra. To illustrate this, KM mixtures created using paper and red dyes for each of the manuscripts can be seen fitted against mean cluster spectra for the different pink areas in Fig 3.47. As can be seen from the two different fitted mixtures, a combination of paper with cochineal carmine fitted well to these pink/red areas, implying again that an insect-based red dye may be the most likely used within these manuscripts.

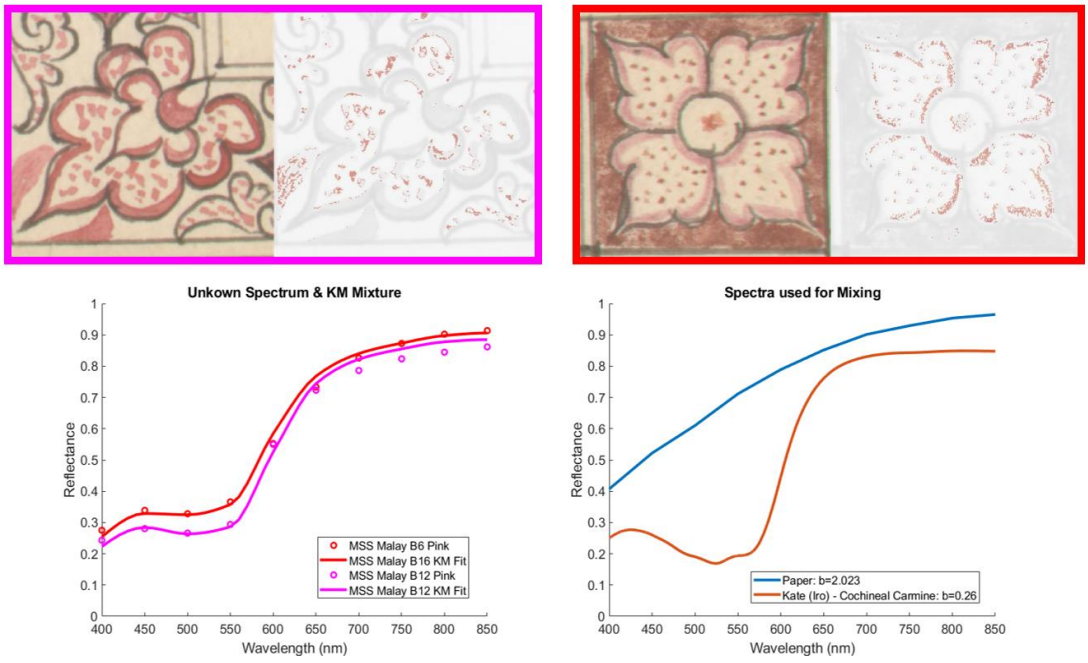

Fig.3.46 Images showing the RGB and true colour cluster maps for a red dye in MSS Malay B6 (Left) and MSS Malay B12 (Right). The bottom shows the KM mixtures for both manuscripts showing cochineal carmine fits well for both manuscripts.

Sulawesi Manuscripts – Add MS 12363

In Add MS 12363, red dye can be seen used for the dark pink coloured circles seen on folio f.37v. To perform analysis of the dye on this page, FORS was used to record the reflectance spectrum of the pink areas, after which KM mixture modelling was then used to create a spectrum to find the best-fitting reference dye. The spectrum extracted from the pink area and the KM fit can be seen in Fig 3.49. As is evident from these plots, the best-fitting dye was found to be cochineal carmine, and by inspecting the absorption features of the unknown red dye (positioned at ≈524nm, ≈561nm) the feature positions agreed with the KM plot where the cochineal carmine reference possessed its own features at 525nm and 560nm, implying use of an insect-based red dye.

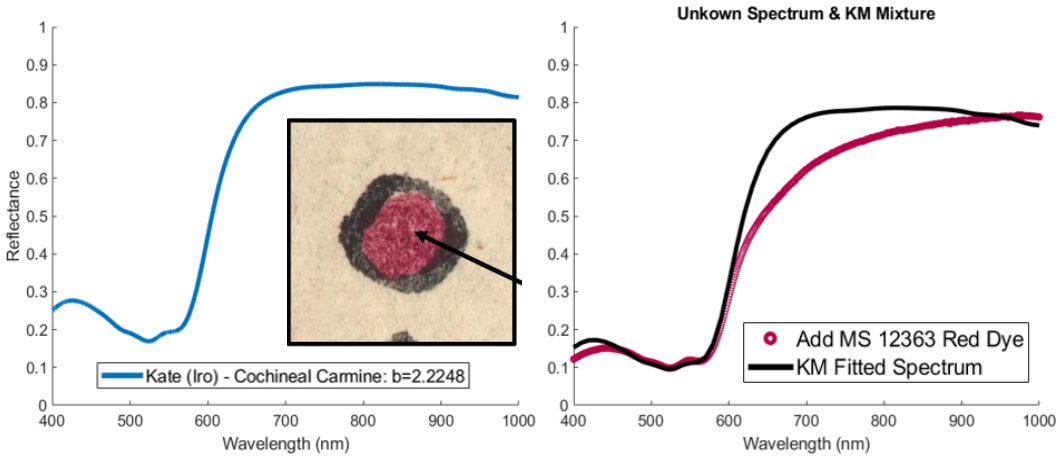

Fig.3.47 KM mixture showing a well-fitting spectrum using cochineal carmine for the red dye areas in Add MS 12363 f37v.

Lombok Manuscripts – Or 15924

Within manuscript Or 15924 there is also suspected to be another red dye. When extracting PRISMS spectra from the outer red regions of the scroll, there is a noticeably weak absorption feature at 450nm which could suggest the presence of a red dye. To investigate the red dye, multiple KM mixtures were created using different dyes and paper, with the best reference dye for the red used on Or 15924 being found to be Madder Lake (see Fig 3.49), however without further study the type of red dye cannot be known for sure.

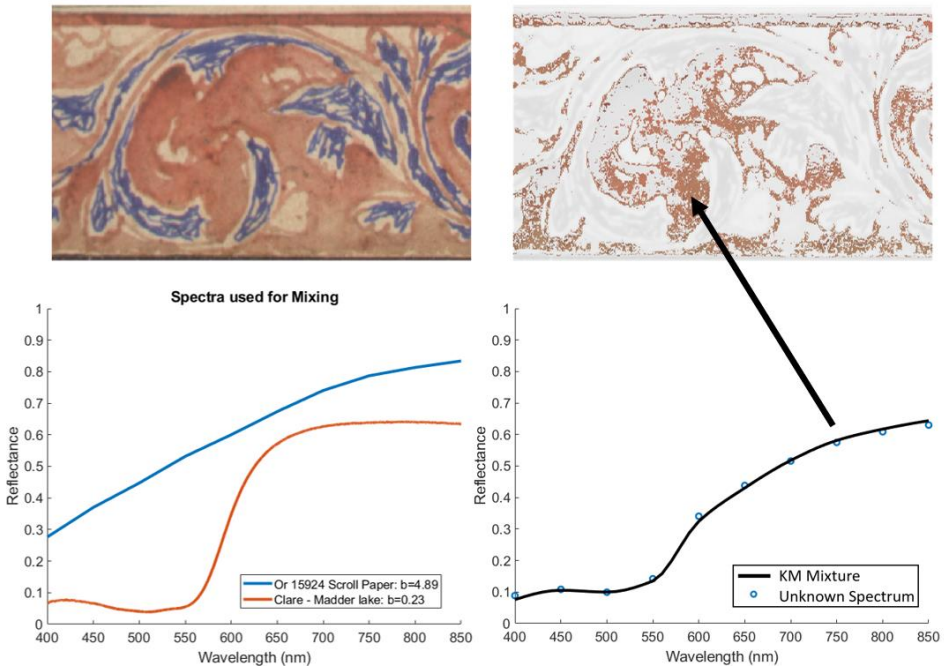

Fig.3.48 Graphs showing well-fitting KM mixtures for a mixture of Madder Lake and paper for the red clusters seen in Or 15924

3.2.3.4 Vermillion + Red Lead

Throughout the Javanese parts of the collection, there are multiple areas in which red lead may be used, however, it is never as a pure pigment and instead tends to be mixed with other reds such as vermillion. Evidence of this can be found in multiple manuscripts including MSS Jav 24 and MSS Jav 68, and MSS Jav 89 where they could be confirmed using XRF and VNIR spectroscopy. In all three manuscripts, the presence of lead, mercury, and sulphur be found in the XRF spectra, as demonstrated in Fig 3.50. There is a chance that these mixtures instead include the use of lead white, but as the white which tends to be used for these manuscripts consists of bone ash, it seems more likely that red lead is used instead. To further demonstrate that a mixture of red lead and vermillion fits well for these areas, a KM mixture using the two as components with the paper was produced with a good fit to a red spectrum extracted from MSS Jav 24 that contained this XRF spectrum (see Fig 3.51). The mixing of red lead with vermillion was often practised in Europe as early as mediaeval times (Coccato et al., 2017) and showed continued use during the 18th century (Mulholland et al., 2017). If the mixture had continued to be used into the 19th century as well, it may be also expected in the manuscripts seen within the collection due to the European influences and use of European paper seen amongst different manuscripts within maritime Southeast Asia at the time.

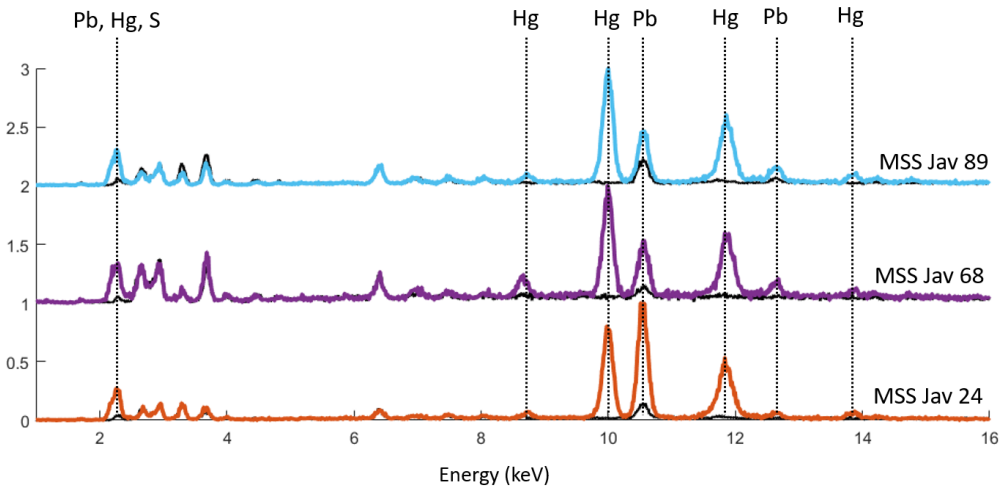

Fig.3.49 XRF spectra for different Javanese manuscripts showing the potential for red lead to be mixed into the vermillion for different red regions.

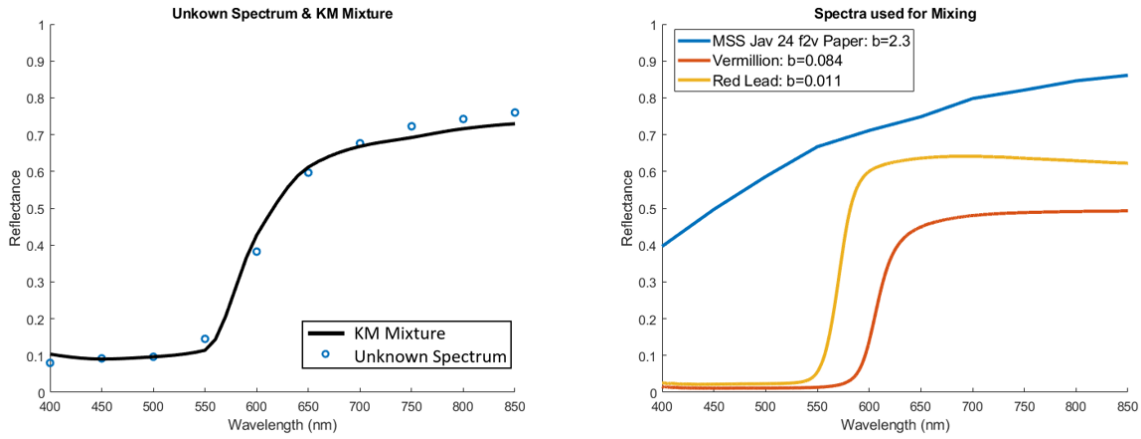

Fig.3.50 KM mixture showing a good fit when paper, vermillion, and red lead are used to fit a spectrum for a red region in manuscript MSS Jav 24 f2v

3.2.3.5 Vermillion + Red Lead + White

Javanese Manuscripts

In a select few Javanese manuscripts there is also evidence to suggest that pink colour is achieved by mixing vermillion and red lead with bone ash white. The use of vermillion and red lead as a mixture has already been discussed in section MSS Jav 24 for instance, XRF spectra recorded for the pink dragon on folio f92v show clear peaks for lead (Pb), mercury (Hg), sulphur (S), calcium (Ca) and phosphorous (P), as shown in Fig 3.52. Furthermore, when creating a virtual spectrum using KM mixture modelling, accurate fits to extracted pink spectra can be made by mixing red led, vermillion, and white (shown in Fig 3.53, where the KM mixture uses a shell white as a substitute for the bone ash white due lack of reliable representative reference spectrum).

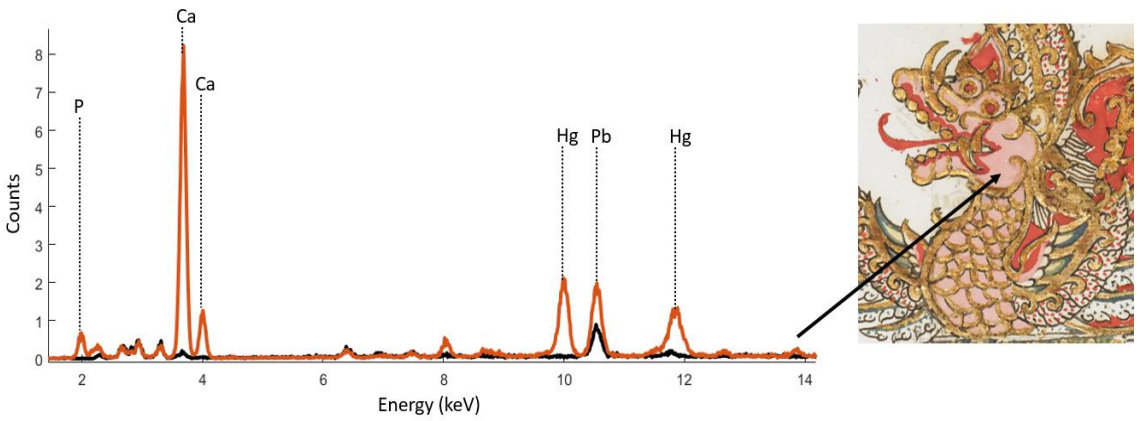

Fig.3.51 XRF spectrum showing the presence of mercury (Hg) and lead (Pb) suggesting the use of vermillion and red lead, mixed with bone ash white (due to the presence of calcium (Ca) and phosphorus(P)).

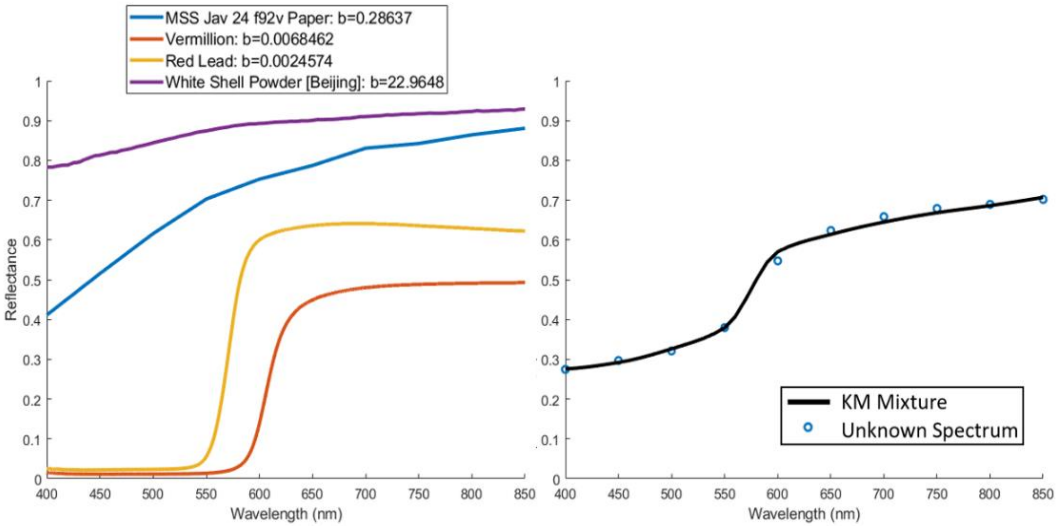

Fig.3.52 KM mixture showing that a good fit for the pink areas on MSS Jav 24 f92v can be achieved using paper, vermillion, red lead, and white (in this case a shell white) spectra.

3.2.3.6 Vermillion + White + Red Dye

Throughout the collection numerous manuscripts use pink colours, however, there is one specific type of pink shade which is suspected to be created as a product of mixing mainly vermillion and white pigments with minor additions of red dyes.

Javanese Manuscripts

These suspected vermillion-dye-based mixtures are primarily found in Javanese manuscripts, where evidence of vermillion and red dye mixed with the bone ash white pigment, detected in 4.1.1, is found to create a pink colour. This can normally be confirmed by using XRF, which shows the presence of mercury, sulphur, calcium, and phosphorus. This mixture can be seen in many different manuscripts including Add MS 12280, 12284, 12285, 12291, 12292 and 12339, as demonstrated by Fig 3.54 which shows the XRF spectra for pink areas in the different manuscripts. Using XRF alone it could be believed that the pink is created purely by mixing the vermillion and bone ash white, however when creating KM to fit with the pink spectra, much better fits can be achieved by also adding additional red dye, such as cochineal lake, as shown through the following example in Fig 3.55 for the pink parts of the flower on Add 12284 f2r.

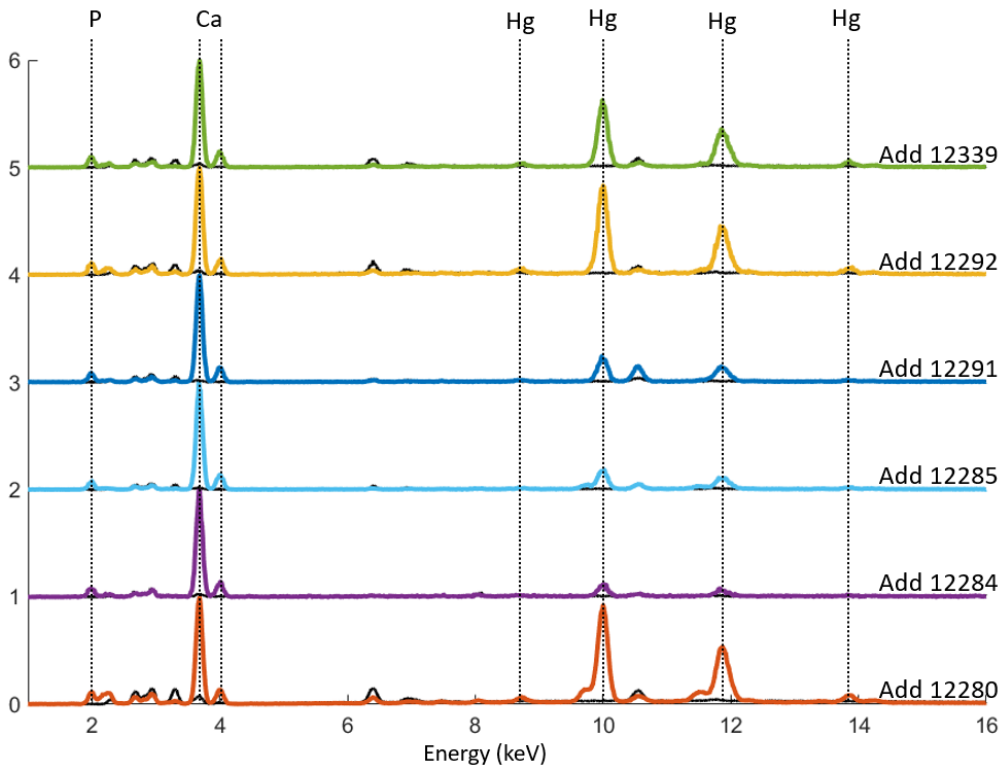

Fig.3.53 XRF spectra for numerous Javanese manuscripts where pink is used, indicating the presence of vermillion and bone ash white due to the presence of mercury (Hg), calcium (Ca), and phosphorus (P).

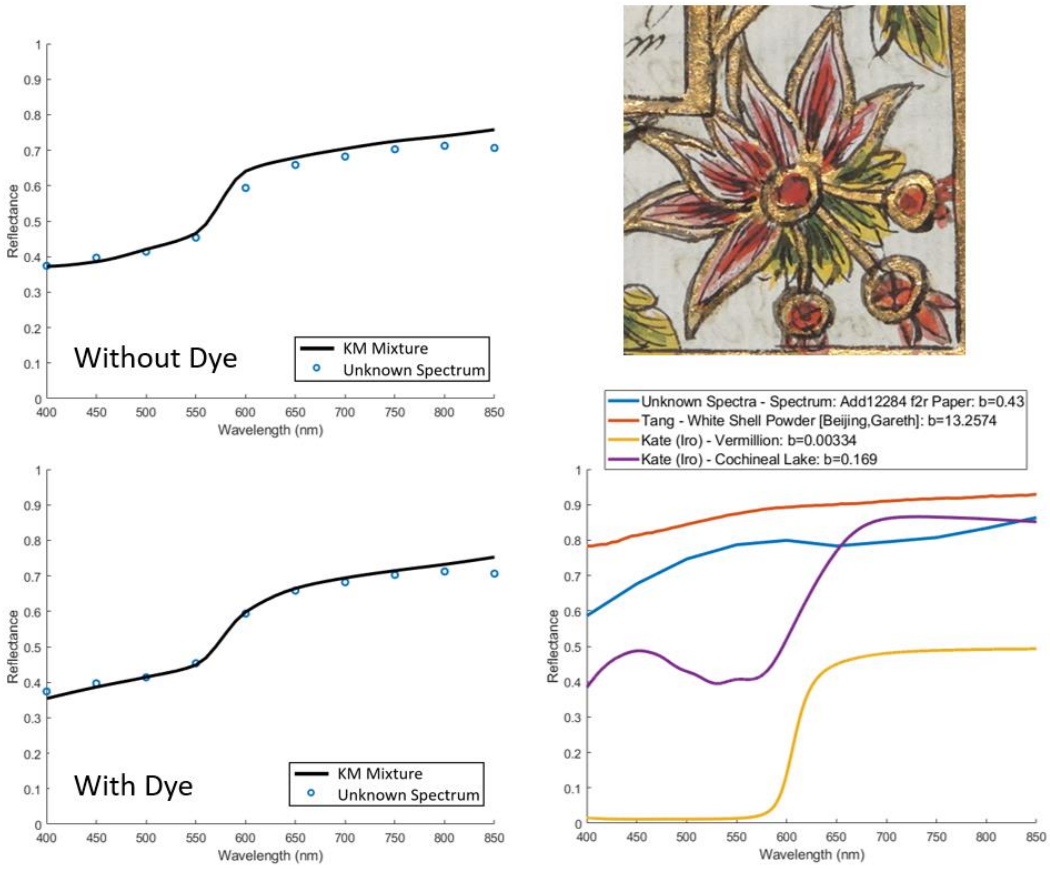

Fig.3.54 KM mixtures showing how well virtual spectra can fit with VNIR reflectance spectra in Add MS 12284 f2r by using vermillion, white, and paper with or without additional cochineal lake.

3.2.3.7 Red Dye + White

Javanese Manuscripts – Or 9333

As was briefly mentioned in a previous chapter, Or 9333 utilises red dyes mixed with white throughout numerous pages. In folio f10v, a mixture can be found to represent the face of a Wayang character, which previously had been predicted to use a mixture of lac lake and ultramarine to create the purple hair/headpiece. The shade of the face however is slightly pinker and possesses a reflectance spectrum which is different to that of the hair, so was analysed further using KM fitting to investigate which materials may have been used to create the pinkish mixture.

It was found, when creating mixtures for the unknown pink, that a mixture of lac lake and a white could be used together to create a virtual spectrum which fitted reasonably well to the main absorption features, as demonstrated in Fig 3.56. With red dye likely being used in the production of the purple pigment seen in this manuscript page, it stands to reason that the same red dye may have been used in other capacities within the same painting, though a confirmation of this is difficult to acquire without further study with FORS.

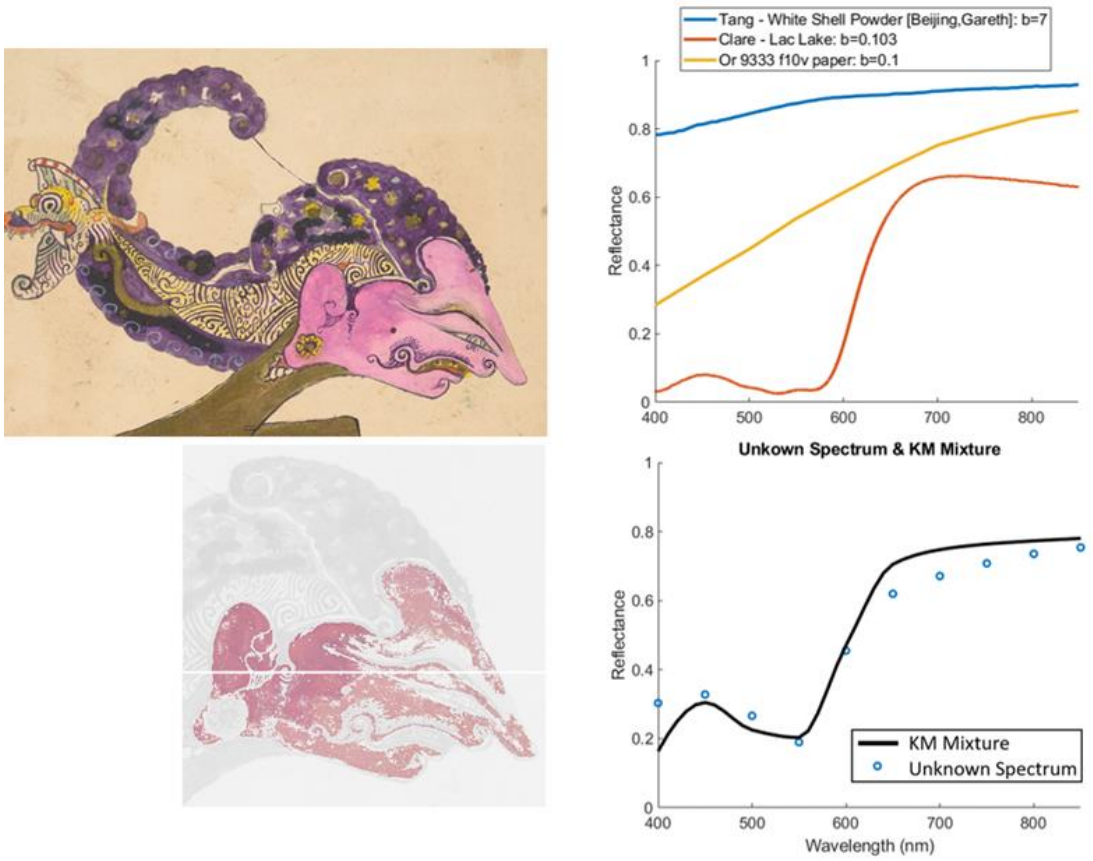

Fig.3.55 KM spectra showing how a mixture of lac lake, paper and white can fit for the unknown pink area seen on the face of the Wayang character in Or 9333 f10v.

Malay Manuscripts – Or 13295

In manuscript Or 13295, a pink colour found on the flower motifs within one half of the double frame located on folio f190v is also predicted to be a mixture of a red dye with white. When analysing the XRF data taken for one of these flowers, sharp energy peaks representing lead can be seen in the spectrum, suggesting that the white pigment used is lead white and therefore must be mixed with a red to achieve a pink hue (see Fig 3.57).

While the pink may be achieved using red lead, extracting VNIR spectra from the flower motifs clearly shows that a red dye fits more closely than any other pigment. For example, in Fig 3.57, a mixture of paper, lead white, and cochineal lake offers a KM model created virtual spectrum which has a very good fit to the true extracted spectrum for the pink. This can be further confirmed using fibre optic reflectance spectroscopy (FORS) which, with its higher spectral resolution, can reveal the clear absorption features indicative of an insect-based red dye (also in Fig 3.57). In the graph shown in Fig 3.58, the best fitting mixture that could be found to fit with the unknown FORS spectrum in the region of 400-850nm was a mixture of lead white and lac lake instead of the previously fitted cochineal, which seems reasonable considering that Lac was used widely over much of Southeast Asia. It is however difficult to be sure of exactly which red dye is used for the mixture, as it is still possible for other dyes, such as cochineal lake which had been fitted to PRISMS spectra, to fit well with the spectrum for the pink flower after being used as components within a simulated KM mixture. Under closer inspection, it was found that the pink flower spectrum possessed two indicative absorption features corresponding with red dye at 532nm and 572nm, which placed this dye close to a reference for lac juice (with features at ≈ 533 and 574nm), therefore suggesting that the pink in this manuscript is achieved with a mixture of lead white and an insect-based red dye, though it cannot be confirmed which red dye is used unless further studies are performed.

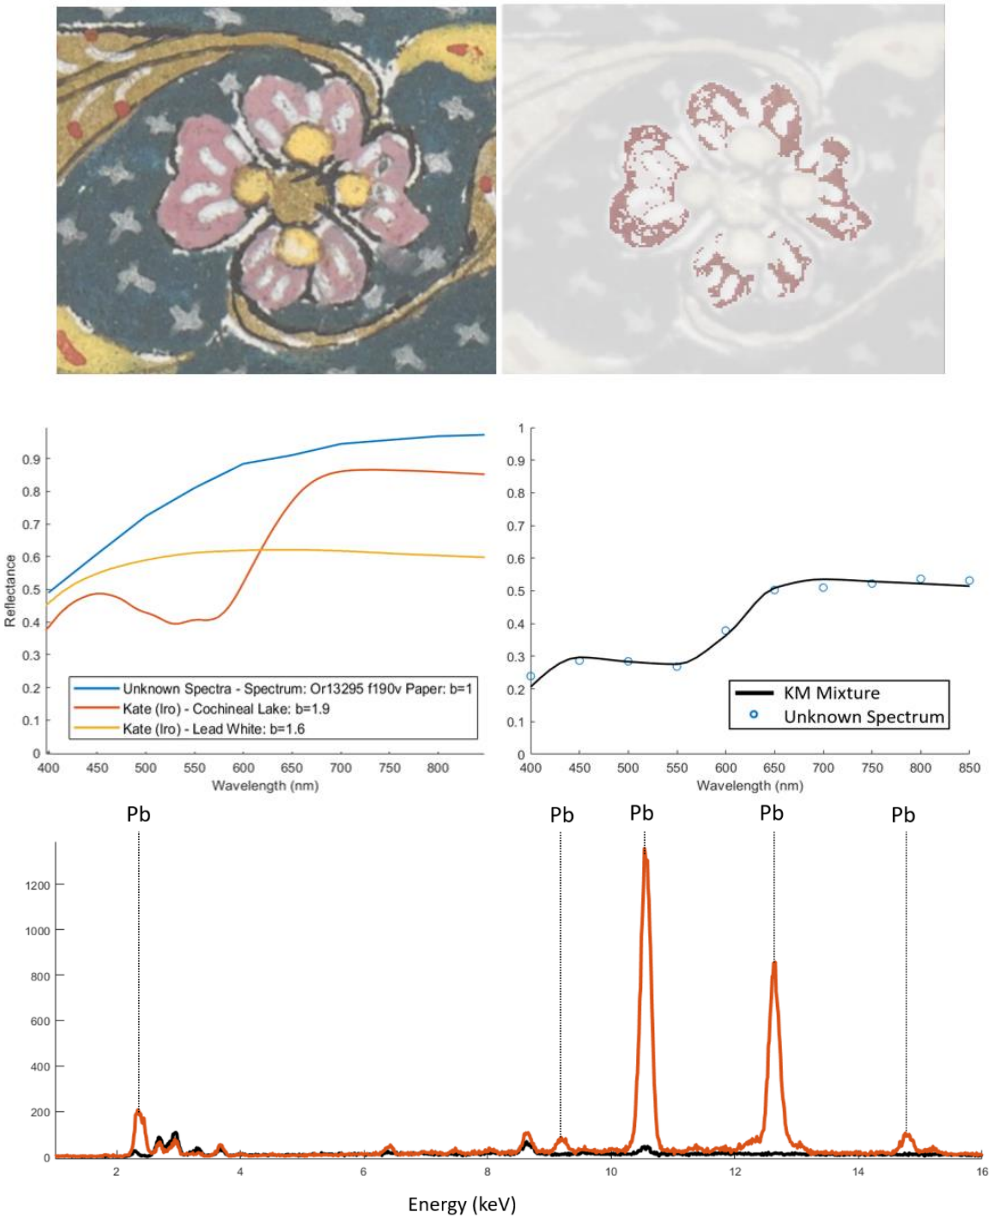

Fig.3.56 Graph showing KM mixture and XRF spectra showing that the pink flower motifs in Or 13295 f190v are likely a mixture of cochineal lake and lead white

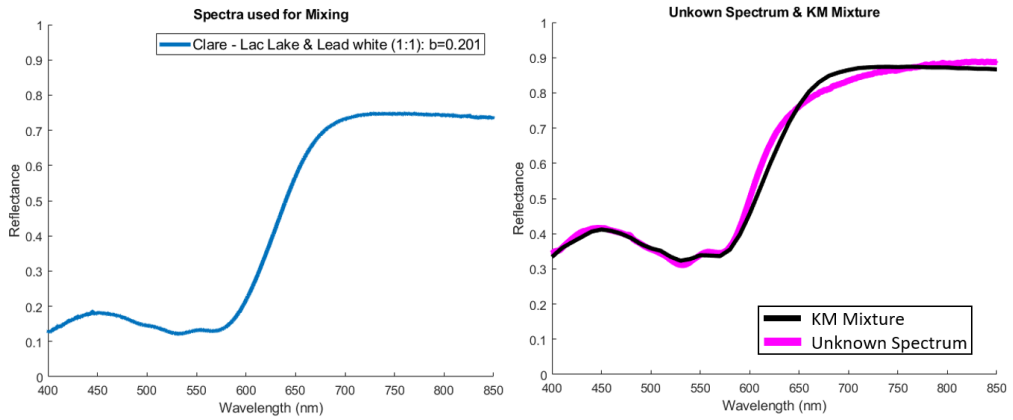

Fig.3.57 KM mixture showing that a mixture of lead white and lac lake fits well to the unknown spectrum in Or 13295.

3.2.3.8 Red Dye + Blue

Javanese Manuscripts - Or 15932 (Dye + Prussian Blue)

As was mentioned briefly in section 3.3.4.3, a mixture of a red dye with blue, likely Prussian blue due to its high absorption, could be detected in Or 15932 (dated to 1814), on folio f.72v. The mixture in question created a dark purple colour, which acted as the background for the genealogical chart included on the manuscript page. To identify the presence of a dye within this mixture, FORS was used to extract a VNIR spectrum, then, by using KM mixture modelling it was found that mixing Prussian blue with a spectrum of the red dye extracted from Or 15932 f.17v resulted in a well-fitting spectrum, as demonstrated in Fig 3.59. For this purple region, absorption features indicative of red dye could be detected at ≈527nm and ≈567nm, which matched closely with a reference for Lac Lake, which possessed features at ≈528nm and ≈568nm, suggesting the use of insect-based red dye.

Javanese Manuscripts – Or 9333 (Dye + Ultramarine Blue)

In the purple areas on Or 9333 f10v, there is also a mixture of blue and dye. Using an extracted spectrum for the purple on f10v to perform a KM fitting showed that a mixture of ultramarine and lac fitted will to the unknown spectrum, as demonstrated in Fig 3.60, however as there is a significant deviation from 650nm onwards, there could also be additional dye of a different source mixed into the purple area. Therefore, if a full identification is to be performed for this dye, further point measurements will be required in future. The use of dye with ultramarine in this manuscript seems reasonable though, as ultramarine is used in abundance over many other illustrations within Or 9333, however, it is important to note that the illumination for these manuscripts may have not all been performed at once and are not created by the same artists, and so the assumption of the same colours being used throughout may not be true.

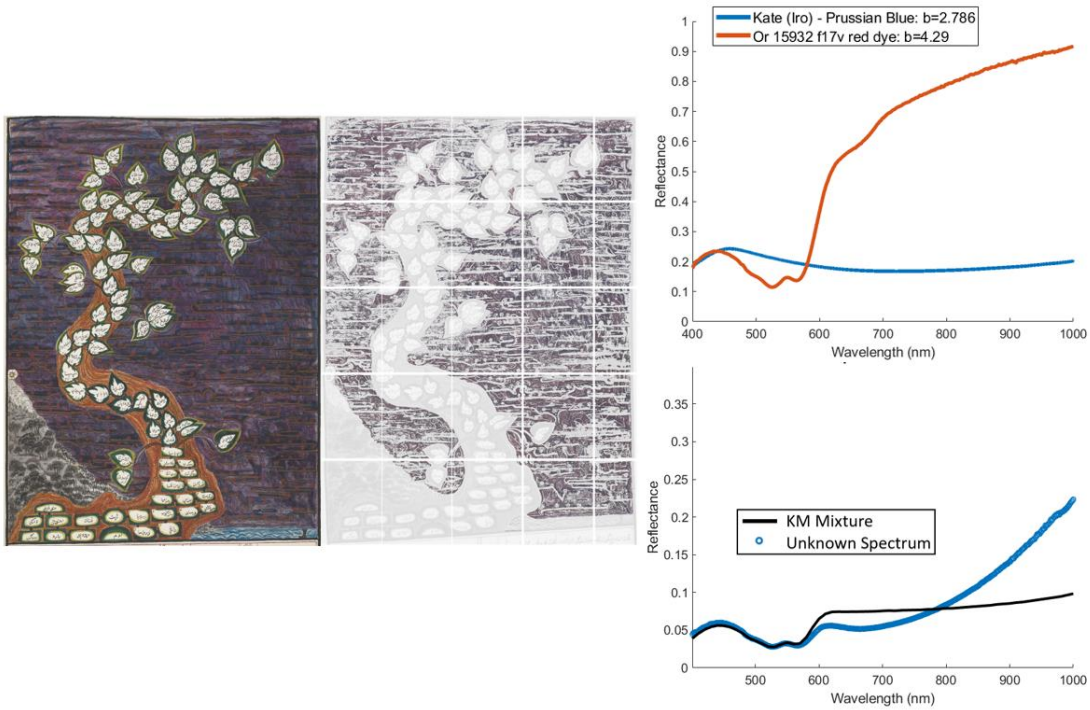

Fig.3.58 Graphs showing that a KM mixture of Prussian blue and red dye (the same dye used elsewhere in the same manuscript) fits well to the unknown purple cluster found in Or 15932 f72v.

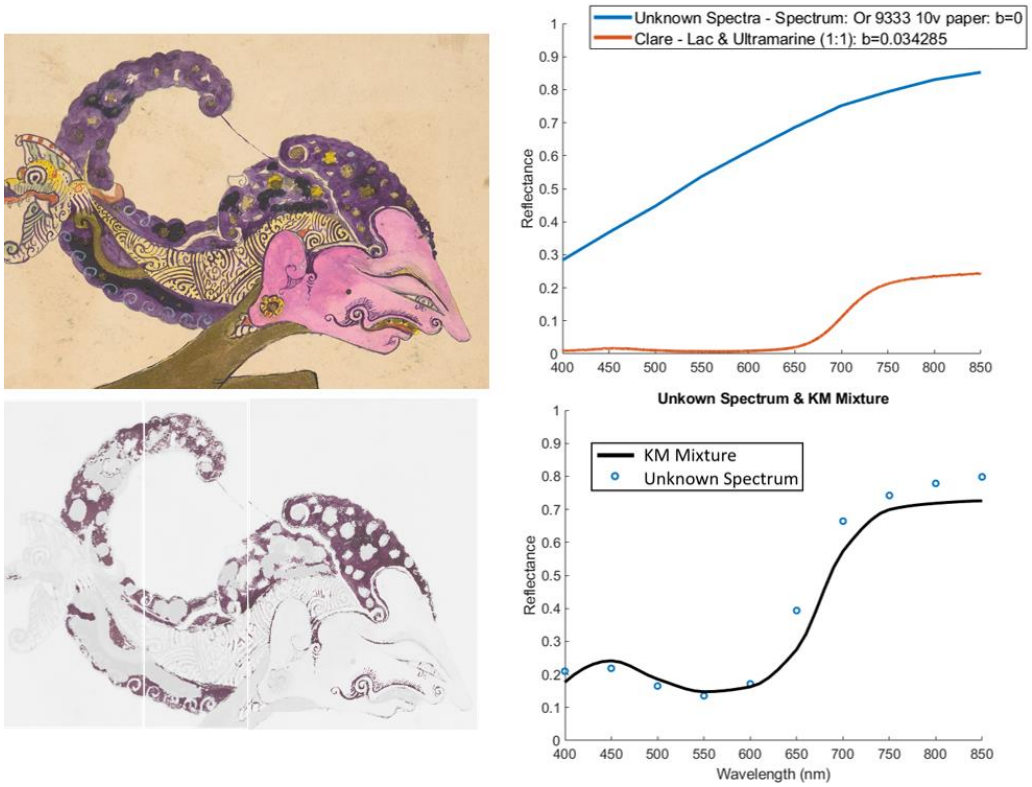

Fig.3.59 Graphs showing that a KM mixture of ultramarine and lac fits well to the unknown purple spectrum found in Or 9333 f10v.

3.2.4 Yellow Pigments & Mixtures

Throughout the collection, yellow pigments are often used in many ways and are found either in the colouring of decorative leaf or flower motifs, as colours in illuminated opening double frames, or the painting of Wayang or human characters, animals and more. Over the many different manuscripts, evidence in favour of four different yellow pigments could be found. These were: orpiment, seen most frequently; followed by yellow ochre; chrome yellow; and an organic yellow which was suspected for one manuscript and confirmed for another to most likely be gamboge.

3.2.4.1 Orpiment

Orpiment is an opaque yellow mineral consisting of Arsenic Sulphide (As_2S_3) and has seen continued use throughout history in many parts of the world as a basic bright yellow pigment, as a component in mixtures to create other colours, or as a paint which can simulate gold. It naturally tends to be found in volcanic environments and therefore can be expected to be found in many areas which have geological volcanic history. In Southeast Asia, orpiment was an important and commonly exchanged trade good often transported to and from locations such as Myanmar (Modern Burma), Malaya, and Thailand from historical mines which produced the mineral as early as the Ming Period (1368-1644), in modern Weishan, China (Kim, 2020). Therefore, as it is well known that many trade routes travelled through these locations into maritime Southeast Asia, it would make sense that orpiment would be used amongst the many different islands if trade had continued through to the 18^{th}-19^{th} centuries. In addition to its production and distribution in China and mainland southeast Asia, orpiment was also a commonly used artistic material in many Central Asian Islamic paintings from the 13^{th}-19^{th} centuries (Knipe et al., 2018) and therefore could have also easily been traded between maritime Southeast Asia and the Islamic world as Islam began to be introduced to local populations prior to the era studied during this thesis. In the modern world, there are documented sources of orpiment within maritime Southeast Asia (Yuningsih et al., 2014), and with most of maritime Southeast Asia having a mostly volcanic origin, it stands to reason that orpiment may have also been locally sourced during the 17^{th}-19^{th} centuries in addition to having been imported. For pigment identification, evidence of orpiment can be found through the combined use of VNIR spectroscopy and XRF. In the VNIR, orpiment typically has a bright yellow colour, very broad spectral features and is opaque, sometimes making it able to be distinguished apart from other yellow pigments. With regards to the XRF, evidence that highly suggests that a yellow pigment is orpiment is through the existence of energy peaks which correspond to both arsenic (As) and sulphur (S), the two main constituent chemical elements of orpiment. However, even with this information the identification of orpiment can be confused with that of pararealgar, another yellow pigment which is also As and S based. Therefore, to be sure that orpiment is used in this collection, Raman spectroscopy was also used.

Javanese Manuscripts

In the Javanese manuscripts within the collection, orpiment was found to be used widely and appeared in almost every manuscript where a yellow pigment was either applied straight onto paper or mixed to create other colours such as orange and green. For many manuscripts, orpiment is the main yellow pigment used, however, there are some circumstances where other alternatives are used as well. Evidence in favour of orpiment was primarily found in the Javanese manuscripts through the identification of its characteristic arsenic (As) and sulphur (S) energy peaks in the XRF, many examples of which can be seen in Fig 3.61. In addition to the XRF however, the presence of orpiment could also be further supported using PRISMS VNIR spectra, where well-fitting KM mixtures using orpiment references could be created for different yellow regions in multiple manuscripts. Some examples of these are shown in Fig 3.61 for Add MS 12297, Add MS 12285, and Add MS 12280.

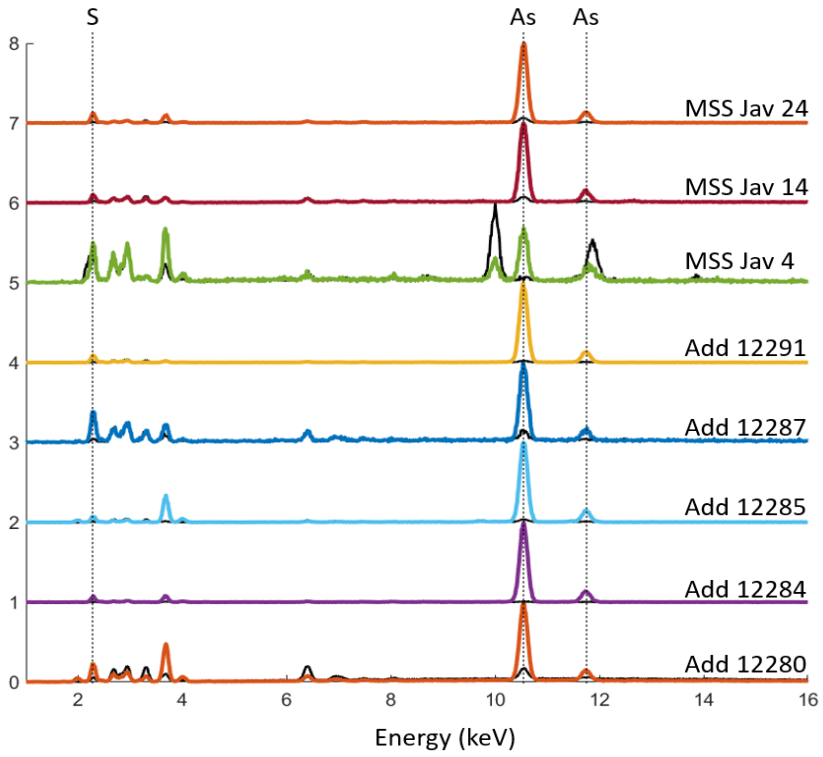

Fig.3.60 XRF spectra showing the presence of arsenic (As) and sulphur (S) indicative of orpiment in numerous Javanese manuscripts.

Furthermore, as can be seen from Fig 3.62, a mixture of orpiment with paper fits well for the yellow regions seen in the different manuscripts, and even good modelling of the transition from orpiment to the paper appears to be achieved at ≈650nm, where the absorption feature of the paper can be seen to dominate quite clearly.

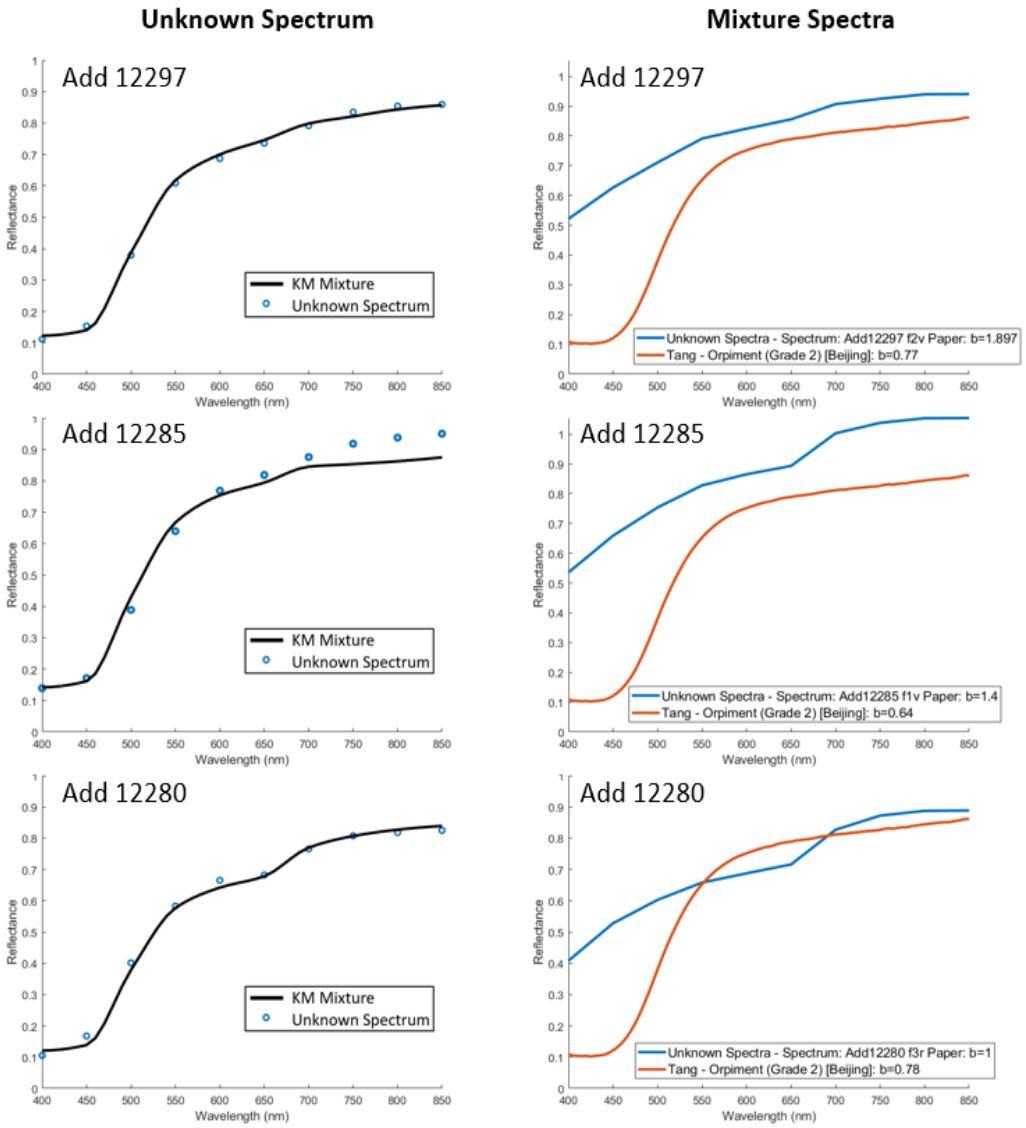

Fig.3.61 KM mixture spectra showing good fits for different yellow regions in Javanese manuscripts when using orpiment and paper.

While the KM mixtures created to fit to orpiment for these manuscripts do fit well, there are other yellow arsenic-based pigments, such as the already mentioned pararealgar, which could also mix well with paper to fit to the unknown spectra. Therefore, for final confirmation of the use of orpiment, Raman spectroscopy was also used to identify characteristic features indicative of orpiment in the same three manuscripts, also shown in Fig 3.63.

With orpiment being confirmed using Raman for multiple yellow areas which had well-fitting KM mixture spectra and As + S peaks in the XRF, it was deemed largely safe to assume that for most Javanese manuscripts, the presence of orpiment could generally be confirmed through the use of XRF and VNIR alone. There is of course a chance that a different arsenic-based yellow could be used in some scenarios, but this is mostly unlikely when considering that all the evidence points toward the continued use of orpiment over most Javanese manuscripts.

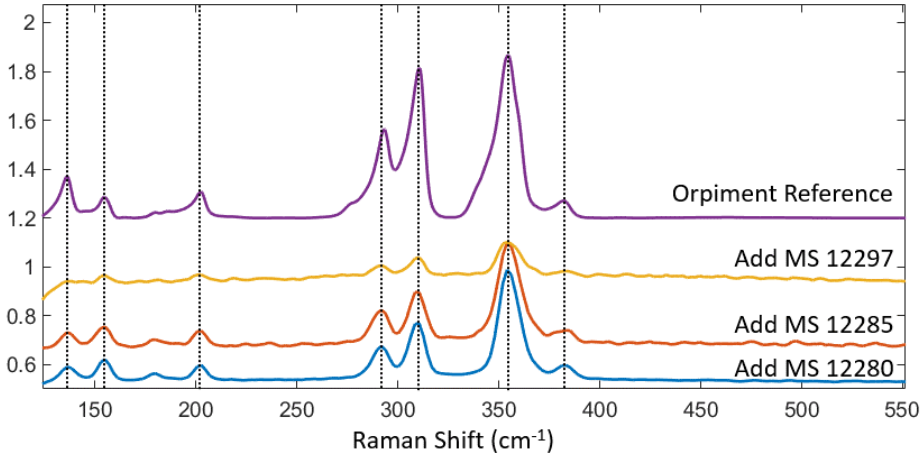

Fig.3.62 Raman spectra indicating the presence of orpiment in three Javanese manuscripts, compared to a reference taken from (RRUFF, 2023b). All spectra were normalised and adjusted for clarity.

Malay Manuscripts

Orpiment was also found in two Malay manuscripts, namely MSS Malay B3 and Or 13295, and again was identified primarily through using XRF, of which the spectra for the manuscripts can be seen in Fig 3.64. For each of the regions, PRISMS spectra extracted from the same locations from which XRF was recorded could be compared against KM-modelled spectra to fit for different orpiment mixtures, the results of which can be seen in Fig 3.65. The existence of orpiment in most Malay manuscripts comes as no surprise as they all have provenance traceable to the Malay peninsula, where, as mentioned previously, there is evidence both for its trade from other locations such as China, but also from local sources (Schafer, 1955). However, as it is also known that orpiment was regularly used in the creation of Islamic manuscripts, even through to the 19th century (Knipe et al., 2018), it would also be reasonable to assume that there would have been continuous trade connections with the Islamic world.

In addition to its existence as a yellow pigment in these two Malay manuscripts, orpiment was also found in another, Add MS 12379 (18th-early 19th CE), where it had been mixed with indigo to achieve a blueish green colour, implying that orpiment was not used only as a pure pigment to achieve yellow colours in Malay peninsula, but also that it could be mixed.

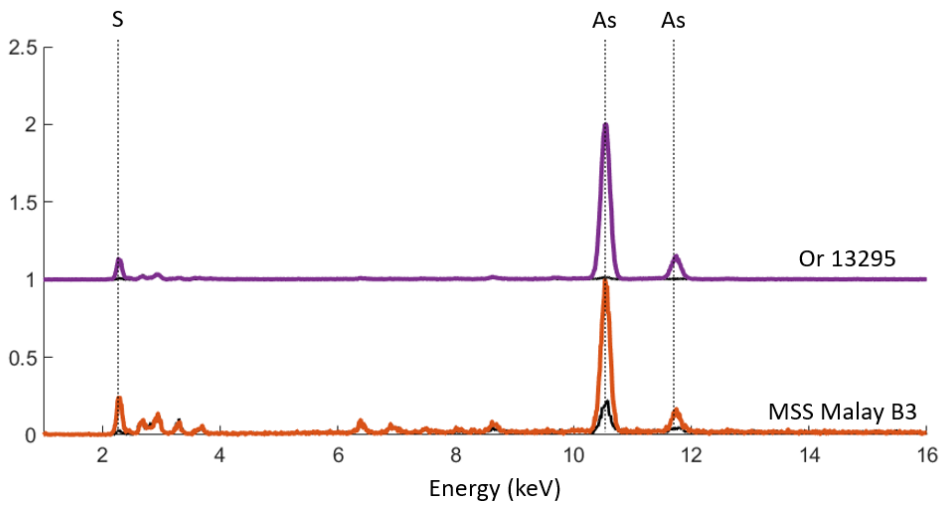

Fig.3.63 XRF spectra indicating the use of orpiment in Or 13295 and MSS Malay B3 due to the presence of indicative arsenic (As) and sulphur (S).

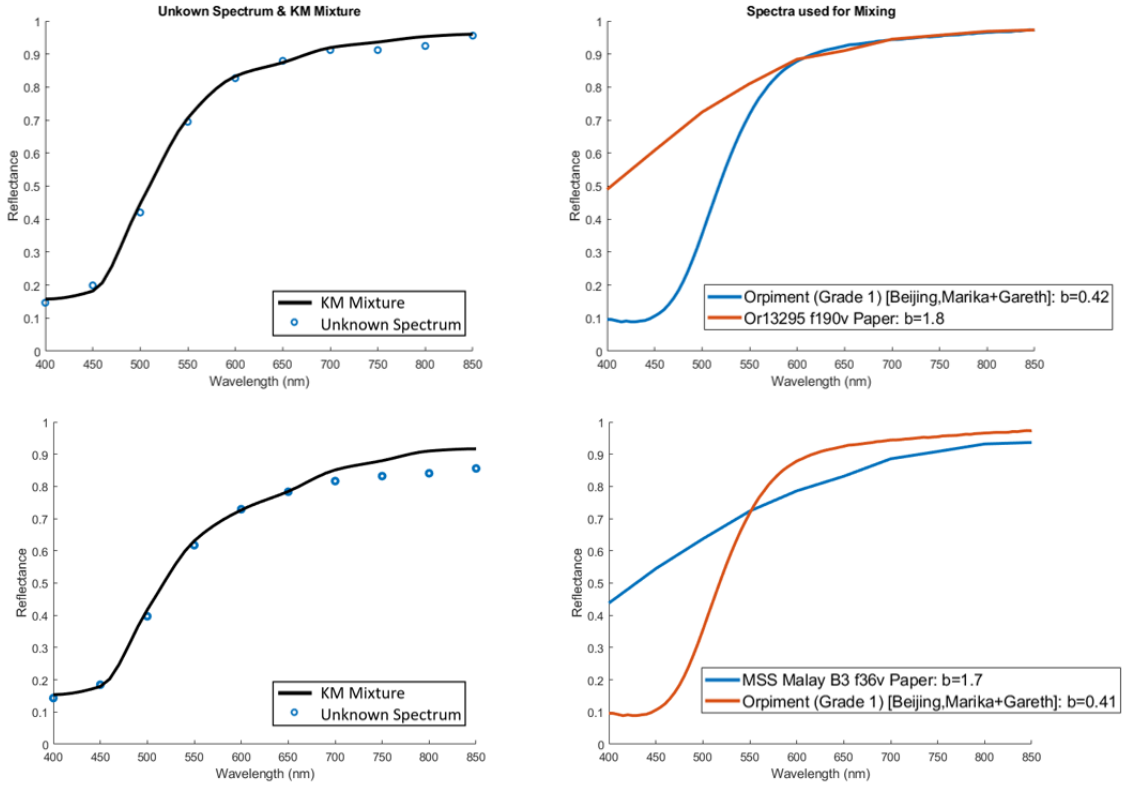

Fig.3.64 KM modelled spectra indicating the use of orpiment due to the existence of well-fitting virtual spectra when using it and paper as mixture components, where the only deviation is towards the NIR.

Sulawesi (Bugis) Manuscripts

In the Bugis manuscripts, two of the three items possessed orpiment where the XRF peaks and KM mixtures of orpiment possessed As and S and fitted the extracted PRISMS spectra respectively. XRF spectra taken from manuscripts Add MS 12363, and Or 8154, can be seen in Fig 3.66.

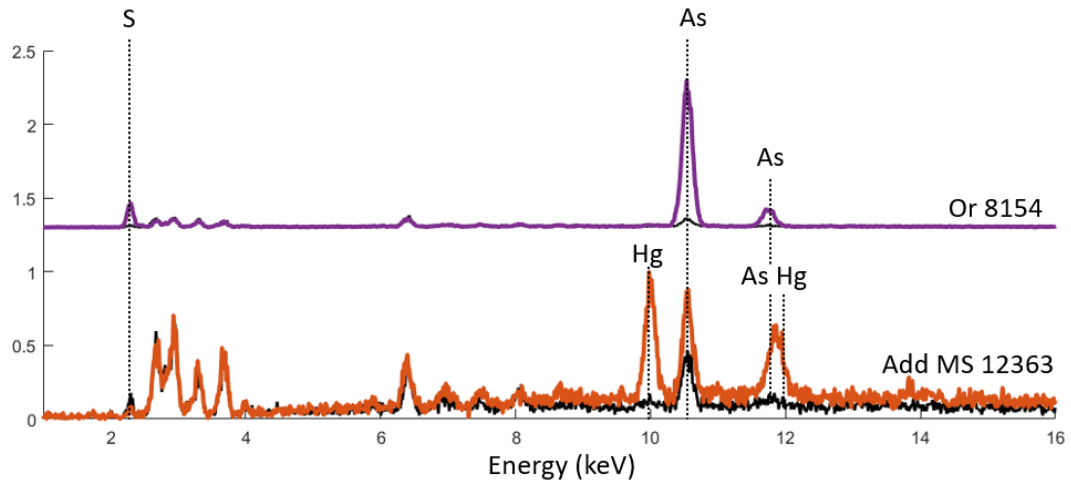

Fig.3.65 XRF spectra indicating the use of orpiment in two Sulawesi manuscripts due to the presence of arsenic (As). In Add MS 12363 there are additional signs of vermillion due to the mercury (Hg) content.

While both XRF plots clearly show the presence of Arsenic, it is strange that for Add MS 12363, the presence of Sulphur in the yellow is lower than that of the paper, furthermore, there are also peaks for mercury in the mixture, which is normally indicative of vermillion for these manuscripts, but the colour of the illuminated area is yellow. It is suspected that vermillion is likely from elsewhere in the manuscript, but the lack of sulphur is strange, however this could be a result of absorption due to the presence of a material layered over the top (which could be the case as the yellow is not a consistent colour).

To provide further evidence and improve the confidence in the yellow pigment being orpiment, KM mixture modelling was used to investigate how well orpiment spectra and paper would fit with the yellow regions in Or 8154 and Add MS 12363. Representative mean cluster spectra, used to avoid the darker and seemingly more degraded regions of the illumination, with fitted KM mixtures for both manuscripts, can be seen in Fig 3.67, where mixtures of orpiment and paper were able to fit the unknowns well.

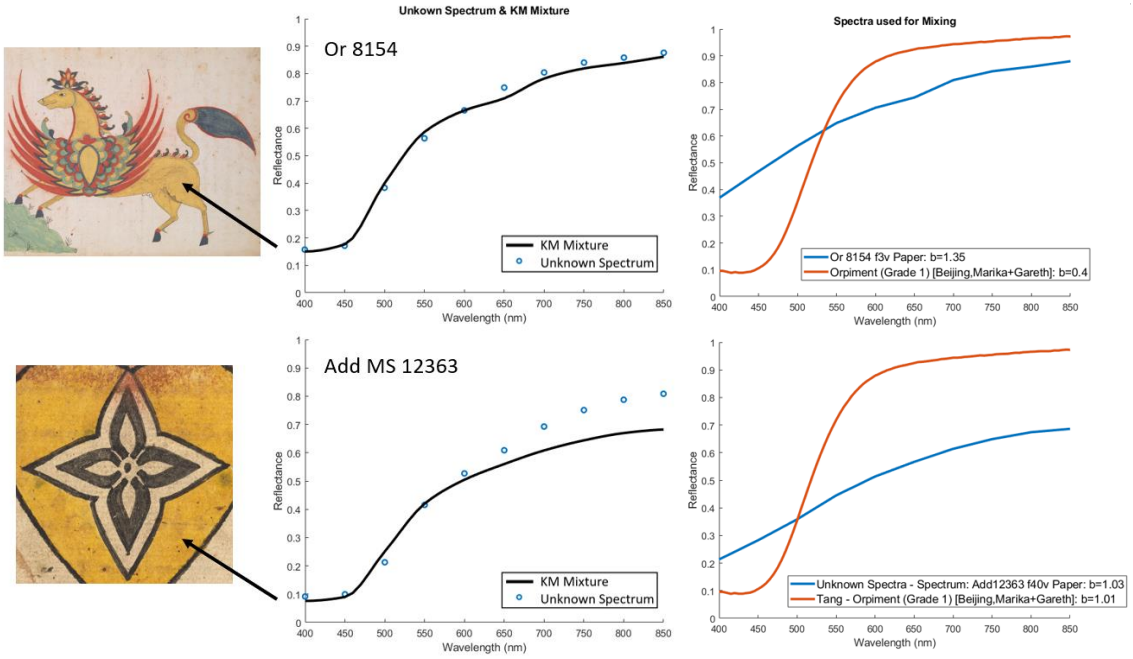

Fig.3.66 KM mixtures showing that a mixture of orpiment and paper fits well with the yellow PRISMS mean cluster spectra taken from Or 8154 and Add MS 12363

Sumatran (Acehnese) Manuscripts

For every Acehnese manuscript where yellow was used, orpiment appears as the main yellow pigment. It can be found in Or 16034, Or 16035, Or 16769, and Or 16915, where the presence of Arsenic (As) and Sulphur (S) XRF peaks exist, as can be seen in Fig 3.68. As can be seen clearly, there is a presence of As and S in all manuscripts, therefore suggesting the presence of orpiment. However, for Or 16034 and Or 16035 there also appear to be additional strong peaks for iron (Fe), which in Or 16034 is stronger in intensity than the peak for arsenic, suggesting that yellow ochre could also potentially be within this mixture as well. If ochre is mixed with the yellow pigment suspected to be orpiment, it is likely that some of the absorption features more indicative of yellow ochre will present themselves in the VNIR spectra. When attempting to fit Kubelka-Munk virtual mixtures to fit to the unknown spectra, however, closely agreeing fits for both Or 16034 and Or 16035 are reached with no ochre being added, as shown in Fig 3.69, it is instead predicted that the iron content is sourced from iron gall ink overleaf.

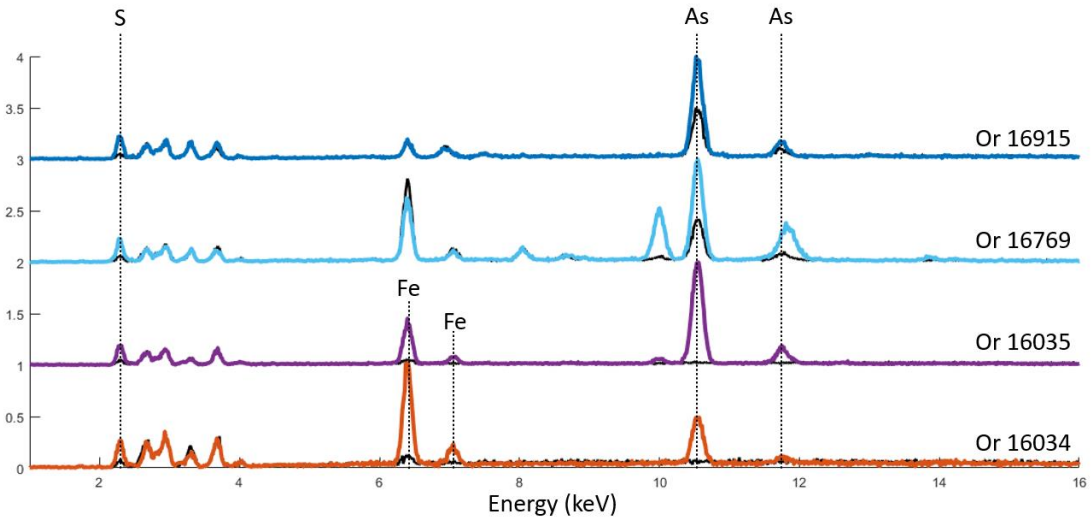

Fig.3.67 XRF spectra indicating the use of orpiment in four Sumatran manuscripts, as shown by the presence of arsenic (As) and sulphur (S)

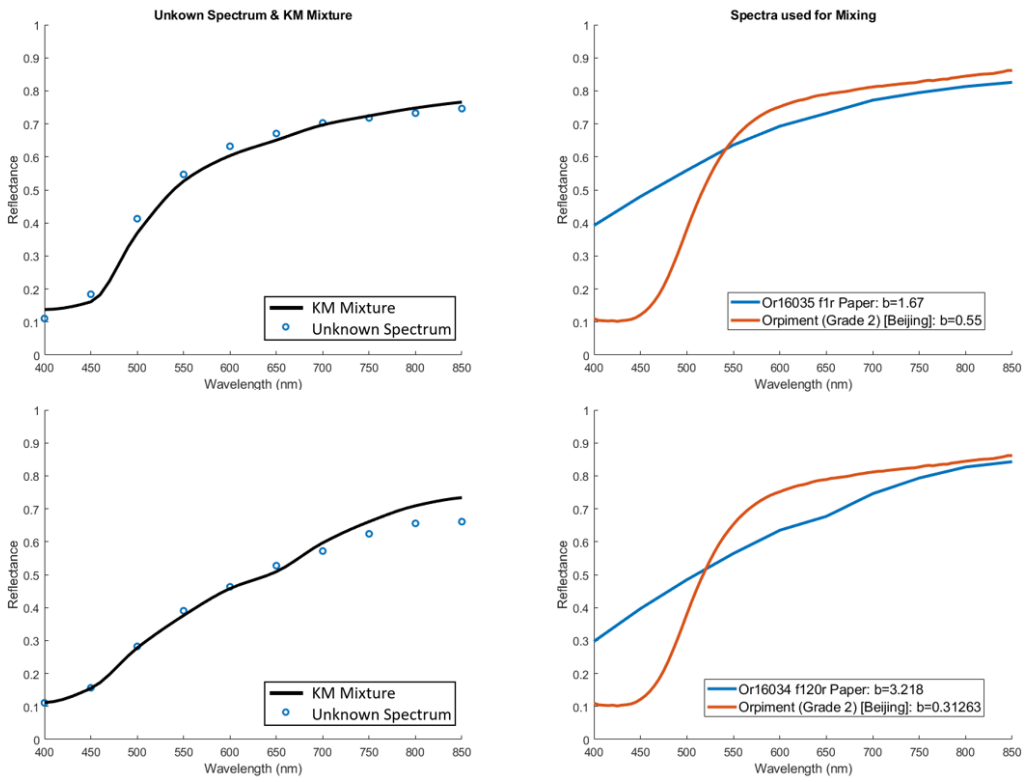

Fig.3.68 Graphs showing how orpiment and paper can be used without an yellow ochre to create well-fitting KM mixture spectra for Or 16034 and Or 16035

3.2.4.2 Yellow Ochre

Yellow ochre is a natural mineral pigment which has seen a long history of use throughout the world and can vary in shade from bright yellow to almost brown with varying levels of opacity. It can be identified by numerous different scientific techniques including FTIR and Raman spectroscopy but was mostly identified in this study using a mix of both XRF and VNIR reflectance spectra. It can be distinguished from many other yellow pigments in the VNIR regime due to changes in absorption over different wavelengths which result in two separate inflection points at ≈500nm and ≈650nm when most other yellows within this collection possess only one S shaped feature between ≈400-500nm.

Javanese Manuscripts – Or 15932

For the Javanese part of the collection, there is evidence of yellow ochre existing on only one manuscript, Or 15932, where it was used as yellow paint to colour the yellow robe of a male character on folio f37r. This can be demonstrated in Fig 3.70, where a KM mixture created using yellow ochre and paper, along with a white pigment to adjust the intensity of the mixture, fitted well with the extracted spectra in the manuscripts, deviating only towards the near-infrared likely due to the transparency of the pigments resulting in a transition from yellow ochre spectral features to that of the paper.

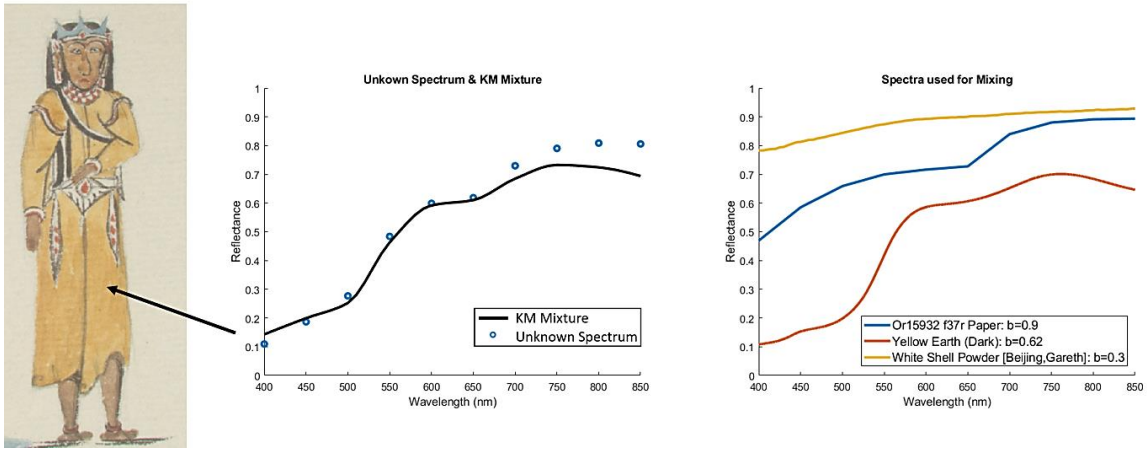

Fig.3.69 Graph showing that yellow earth (ochre) can be used with a white pigment and paper to create a KM mixture spectrum which fits closely with the yellow seen in Or 15932 f37r

Sulawesi (Bugis) Manuscripts – Add MS 12372

There is evidence which also shows the existence of yellow ochre in a single Bugis manuscript, Add MS 12372, however, it can be difficult to be sure of this as in the VNIR, most of the indicative absorption features which normally make ochre straightforward to identify in the VNIR have been lost due to it being mixed with other materials. For example, a virtual KM mixture using only yellow ochre and the paper deviates from the unknown yellow spectrum from ≈650nm, where the

extracted PRISMS spectrum then does not have the other absorption feature typical of yellow ochre found after ≈750nm (see Fig 3.71). However, to gain a better understanding of what the mixture may be, XRF taken from the same region can be shown to possess peaks typical of yellow ochre, as both iron (Fe) and manganese (Mn) can be detected to have counts much higher than the paper. However, there are also additional peaks for lead (Pb), mercury (Hg), and sulphur (S), suggesting that vermillion or a lead pigment may have also been added (also demonstrated in Fig 3.71). This could be the source of why there is disagreement toward the NIR parts of the spectrum.

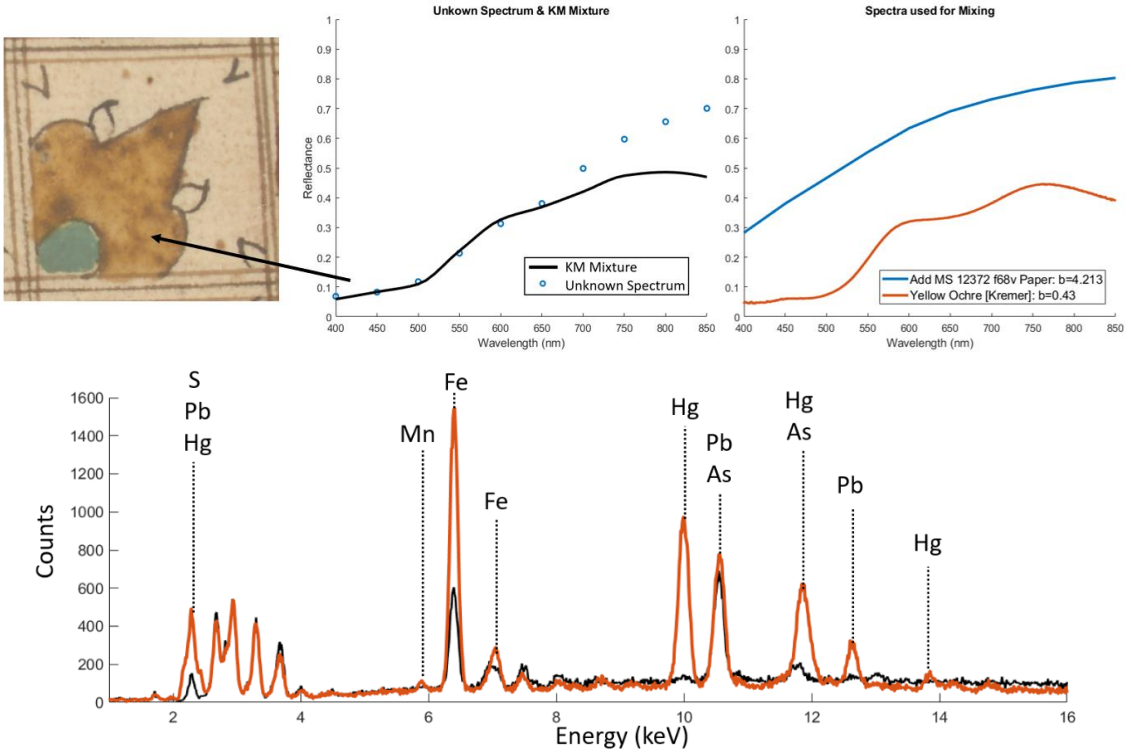

Fig.3.70 KM mixture and XRF spectra indicating the presence of yellow ochre used in the yellow motifs seen in Add MS 12372 f68v.

As the yellow is quite dark and possesses an orange hue, it could also be possible that in addition to yellow ochre, vermillion, perhaps mixed with red lead, may also exist as a component within the mixture. To test this a KM mixture was created again to see if this mixture would be a very good fit for the unknown spectrum, the result of which is demonstrated in Fig 3.72.

Interestingly, the result from adding vermillion and red lead to the mixture results in an even worse fit than just yellow ochre, as the vermillion and red lead are causing the absorption feature of the ochre at ≈600nm-650nm to be weaker. Because of this, it was assumed that yellow ochre is the main yellow pigment used here, however to be understand what other pigments may be present further study would be required.

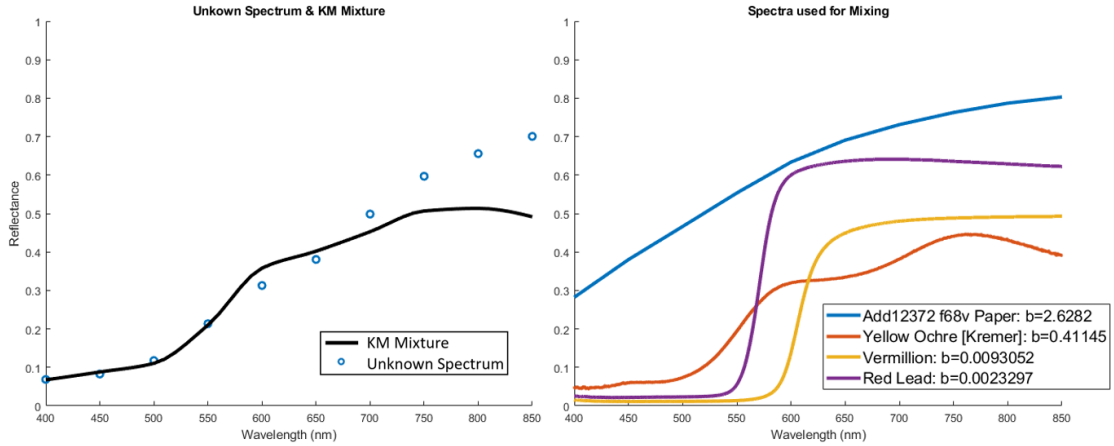

Fig.3.71 Graphs showing the KM mixture creating a worse fitting mixture when vermillion and red lead are included for the yellow motifs in Add MS 12372 f68v.

3.2.4.3 Chrome Yellow

Lastly, chrome yellow was another yellow pigment found within the collection, though rarely and only in the latest manuscripts. Chrome yellow is a yellow-orange pigment that was discovered in the late 18th century but was not so widely used due to its tendency to fade due to sunlight or blacken over time due to exposure to sulphur from the air. It can be detected using XRF due to its inorganic nature, where the chemical formula for chrome yellow is $PbCrO_4$. It is created using lead nitrate ($Pb(NO_3)_2$) and potassium chromate (K_2CrO_4), implying that in some circumstances, impurities such as potassium (K) may be detected in addition to lead (Pb) and chromium (Cr) when using XRF. In addition to XRF, chrome yellow can also be identified using FTIR and Raman spectroscopy. Chrome yellow was not widely available until at least the early 19th century, when deposits of chromium-containing mineral deposits were known to have been mined in France and the United Kingdom in 1818 and 1820 respectively (Feller, 1986). By the 19th century, the British Empire had a significant enough presence that chrome yellow could have been traded from the United Kingdom to maritime Southeast Asia, especially considering that many manuscripts in modern British collections have a history of belonging to figures such as Sir Thomas Raffles and Colonel Colin Mackenzie.

Javanese Manuscripts

Of the Javanese manuscripts, chrome yellow was only found in Or 15026, where it appears as a pigment both for illuminating Wayang characters and for the traditional opening double frames. Evidence in favour of the existence of chrome yellow was found through the identification of lead and chromium peaks in the XRF spectra, as demonstrated in Fig 3.73. Furthermore, additional KM mixtures using chrome yellow were also deemed to possess a close fit to extracted PRISMS spectra, also shown in Fig 3.73.

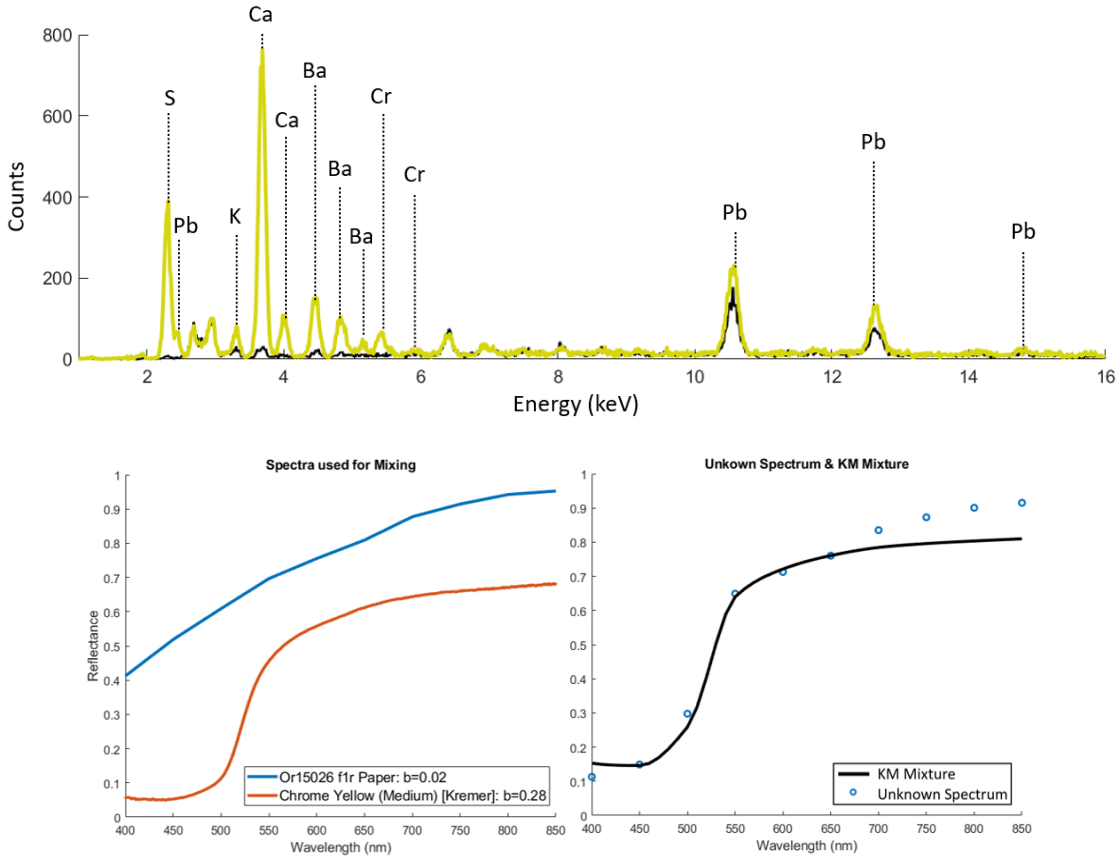

Fig.3.72 (Bottom) KM mixture spectrum indicating that a mixture of chrome yellow and paper fits well to the yellow regions on Or 15026 f1r. This can be further indicated by XRF (top) which shows lead (Pb) and chromium (Cr).

In this manuscript, chrome yellow appears to be mixed with lots of other materials, as XRF shows the presence of stronger-than-expected peaks for lead, calcium, barium, and sulphur. This could suggest that the yellow used is a mixture of several yellow pigments, such as lemon yellow or barium yellow, or that it could simply be mixed with a barium white or lead white mixture.

Malay Manuscripts

Chrome yellow was also found a single time in the Malay part of the collection, in manuscript Or 15227, where it was used in the illumination of many decorative areas including the rubrication seen on every 13th page of the manuscript and in the traditional opening double frames. Again, its use was mainly determined by analysing the XRF spectrum extracted from a yellow region, this time on the border of the frame where both lead and chromium was detected, though with PRISMS reflectance spectra, it could be further confirmed. The different spectra associated with the use of chrome yellow in Or 15227 can be seen in Fig 3.74.

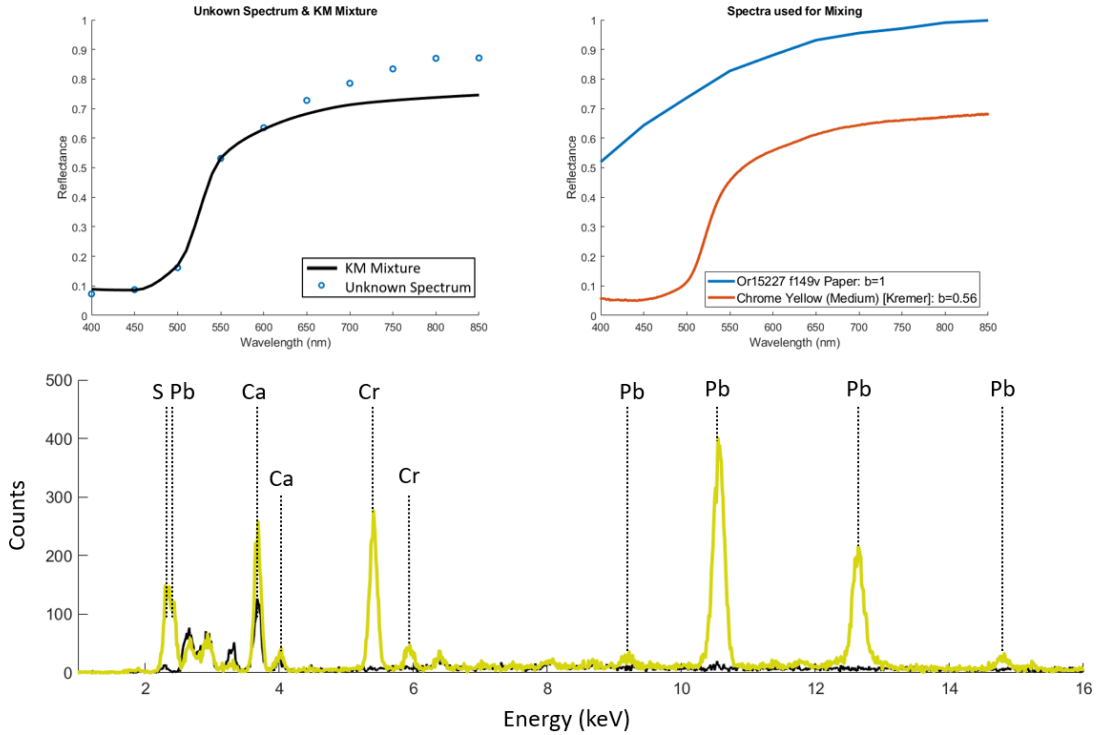

Fig.3.73 (Top) KM mixture spectrum indicating that a mixture of chrome yellow and paper fits well to the yellow regions on Or 15227 f1r. This can be confirmed by XRF (bottom) which shows a strong presence of both lead (Pb) and chromium (Cr).

Unlike in the example of its use in the Javanese manuscript Or 15026, the chrome yellow used in Or 15227 does not appear to be as mixed with as many other materials, as illustrated in the XRF and ER-FTIR spectra. Instead, the XRF spectrum seems to possess only chrome yellow, perhaps with a lead white or calcium carbonate-based white as a mixing component to lighten the yellow, as the main peaks found are for Pb, Cr and Ca only.

3.2.4.4 Yellow + Red Orange Mixtures

Aside from the use of the main yellow pigments, there are also applications of orange in numerous manuscripts which normally exist by mixing red and yellow pigments together instead of using a more naturally occurring pigment. In all the scenarios where orange is found throughout the collection, the mixture used appears to always be a combination of the red and yellow used over the same manuscript. For example, in MSS Jav 24, the orange tends to be a mixture of orpiment with a vermillion/red lead mixture, and in others such as Add MS 12285 it is just a mixture of the vermillion and orpiment already detected elsewhere on the manuscript (as demonstrated in Fig 3.75). This can even be further confirmed by Raman spectroscopy for the orange dragon in Add MS 12285, where features indicative of both vermillion (cinnabar) and orpiment could be found, as demonstrated in Fig 3.76. The orange mixture which is used most is a mixture of orpiment with vermillion, where it is often seen implemented in the motifs surrounding opening double frames

around several manuscripts, the second most common is the red lead and vermillion with orpiment mixture. These orange mixtures can also sometimes be seen mixed with black to make a dark orange or brown colour.

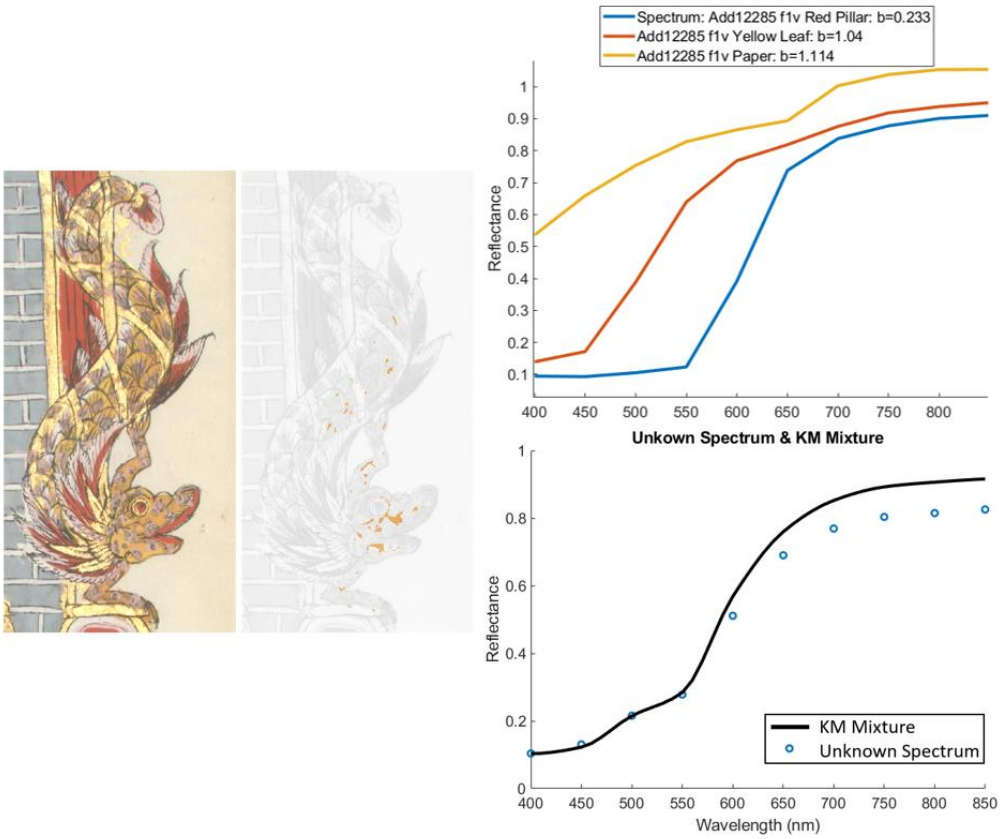

Fig.3.74 KM mixture showing that the same red and yellow (vermillion and orpiment) used throughout Add MS 12285 produces a well-fitting spectrum for the orange regions on the opening double frames.

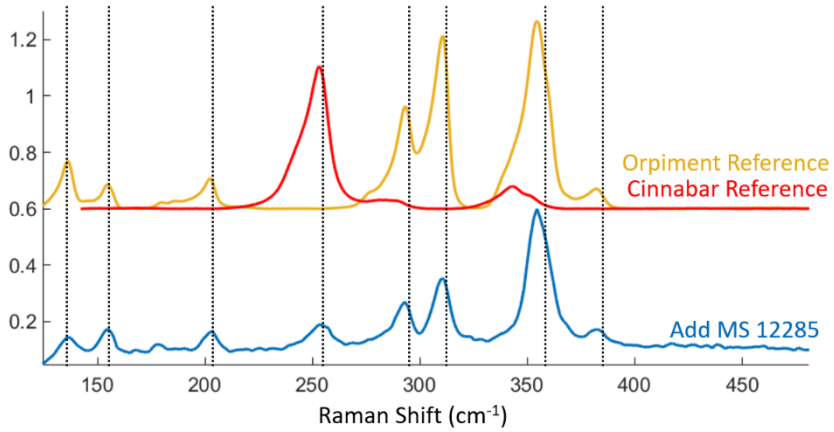

Fig 3.75 Raman spectra taken from (RRUFF, 2023a, 2023b) confirming the presence of both vermillion (in red) and orpiment (in yellow) for the orange areas on Add MS 12285 (in blue).

3.2.4.5 Organic Yellow

There are other yellow regions present on two manuscripts which are suspected to be organic yellow. While there are multiple organic materials used throughout Asia, it is suspected that gamboge is the most likely yellow pigment used for these manuscripts, however conclusive evidence for this was difficult to gather.

Javanese Manuscripts – Or 15932 & MSS Jav 67

The evidence in question appears in two separate manuscripts, MSS Jav 67, where it was detected as part of a green mixture on folio f1r (covered later), and Or 15932, where it was used as a yellow colour on folio f72r. Both manuscripts are similar in design and artistic style to one another, with a seemingly more western influence seen in the many illustrations and illumination. It is therefore highly likely that they are connected in some way.

For the detection of the organic yellow in Or 15932, an XRF spectrum taken for the centre of a yellow and blue star presented very weak energy peaks corresponding to calcium (Ca) and iron (Fe) which were stronger than the paper substrate, potentially suggesting either the use of yellow ochre or an organic yellow with iron as an impurity. However, to check if the material used is yellow ochre, a VNIR spectrum was also extracted to investigate how well either an organic yellow or ochre would fit via the use of a KM mixture.

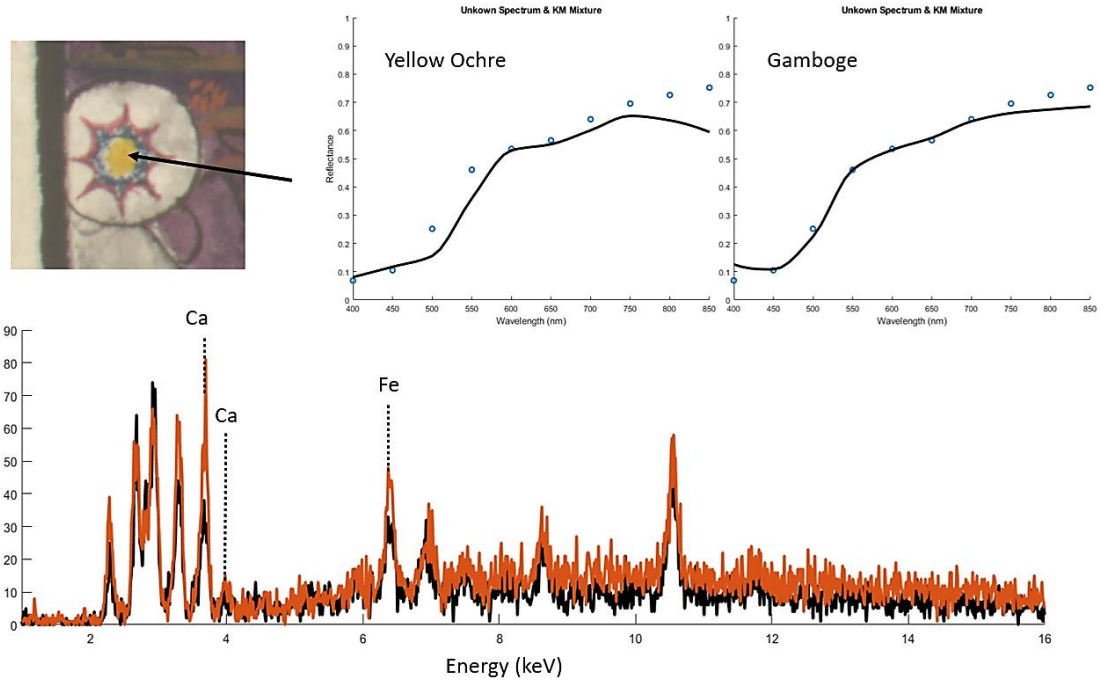

*Fig.**3.76** (top) KM mixture spectra showing that gamboge fits better than yellow ochre for the yellow circle on Or 15932 f72v, even though there is iron content detected in the XRF (bottom)*

When using the paper and either yellow ochre, or gamboge, as components in a virtual mixture, the best matching KM fit was achieved with gamboge instead of yellow ochre. The associated

spectra for both the KM mixing and the XRF can be seen demonstrated in Fig 3.77. Attempts to identify this pigment using other complementary techniques such as SWIR FORS, ER-FTIR and Raman spectroscopy provided no clear features which were able to correspond with any organic yellows. However, with gamboge being detected in MSS Jav 67 (see 3.4.6.3), it was likely that gamboge may have been the same yellow used in this manuscript as well, however, without further study, this cannot be known for sure.

3.2.5 Green Pigments & Mixtures

The colour green holds great importance in many different cultures with influence in maritime Southeast Asia. For example, in modern Islam, the colour green is held with great respect and has many traditional ties with the religion, where it is associated with important religious figures, such as Al-Khidir, and even paradise itself, though this is not to say that the use of green in maritime Southeast Asian manuscripts in the 18th-19th centuries is directly linked to the religion, as there are many different cultures with a history of using green in artistic works. Although, in Islamic manuscripts both copper greens and mixtures consisting of indigo and orpiment were used regularly from the 15th century (Knipe et al., 2018), suggesting that they may have also found use in maritime Southeast Asian manuscripts. In many circumstances, green can generally consist of a mixture of two colours instead of being a naturally occurring colour. The colours which tend to be used for this mixture are typically yellow and blue, though there are some variations where darker brown-green hues can be produced with orange colours in place of yellows. Some typical examples of this which may be expected to be found in maritime Southeast Asia, in addition to the previously mentioned indigo and orpiment, can also include “Hooker’s green”, a mixture of Prussian blue with gamboge, or “grass green” a common mixture of indigo and gamboge found in Chinese paintings (Mccarthy & Giaccai, 2021).

3.2.5.1 Copper Greens (Malachite & Mixtures)

Copper-based green pigments have been used in Asia and Southeast Asia for many centuries, and they were found many times throughout this collection. It is important to mention a caveat of this study however, insomuch that the copper greens in this collection often could not be fully characterised. Therefore, malachite is often used as a “best fit” constituent for KM mixtures, but more as a placeholder for estimating the use of green copper pigments in numerous manuscripts instead of performing full pigment ID (hence why this section remains named “Copper Greens”).

Javanese Manuscripts

Within the Javanese part of the collection, there is evidence in favour of copper greens being used on four Javanese manuscripts: MSS Jav 28, MSS Jav 68, and Or 9333. This can be seen with the

use of XRF in all the manuscripts (Fig 3.78). There is also evidence in favour of a copper green being used in MSS Jav 17, however, no point measurement data apart from PRISMS was recorded for this location.

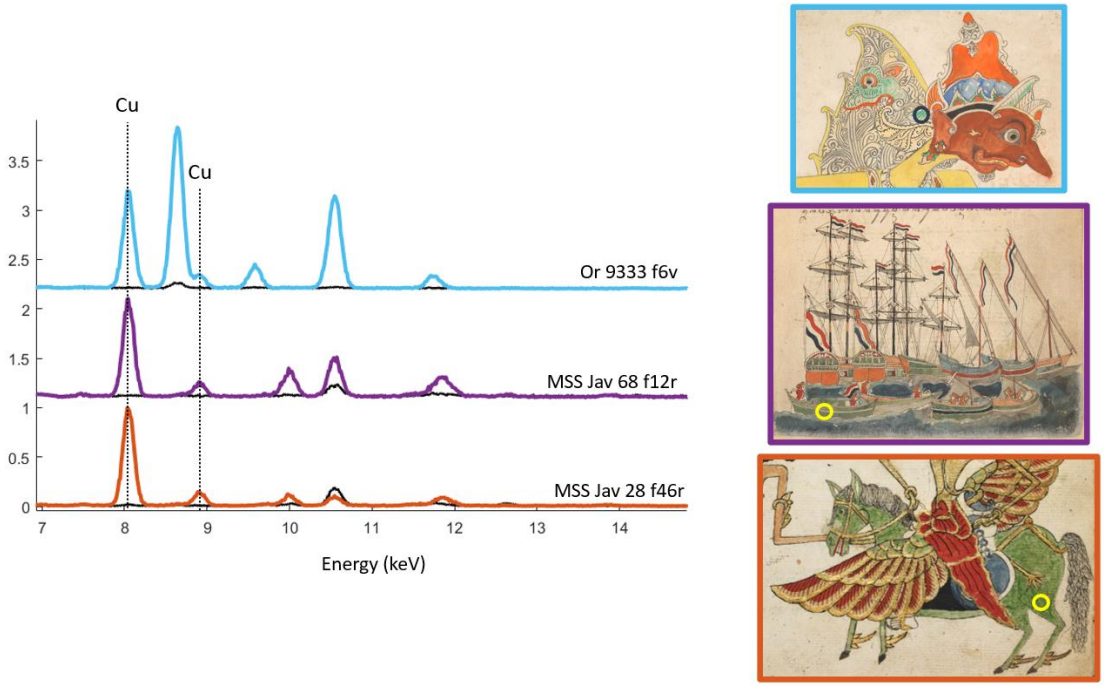

Fig.3.77 XRF spectra indicating the use of copper greens in three Javanese manuscripts.

In all these XRF spectra, there is an increased presence of copper (Cu) which can be found in the different green areas of the manuscripts, where counts are higher than in the paper. For all these manuscripts, VNIR spectra fitted well with mixtures that included malachite as a placeholder, therefore when analysing the green in MSS Jav 17, where XRF was not recorded, a KM mixture using paper, and malachite, with a white pigment to adjust the intensity of the simulated spectrum, was used for the identification. This mixture was found to fit very closely with the unknown spectrum for the green area of the Wayang character on folio f169r, as demonstrated in Fig 3.79, and therefore suggested that the green pigment used on this MSS Jav 17 was also copper based.

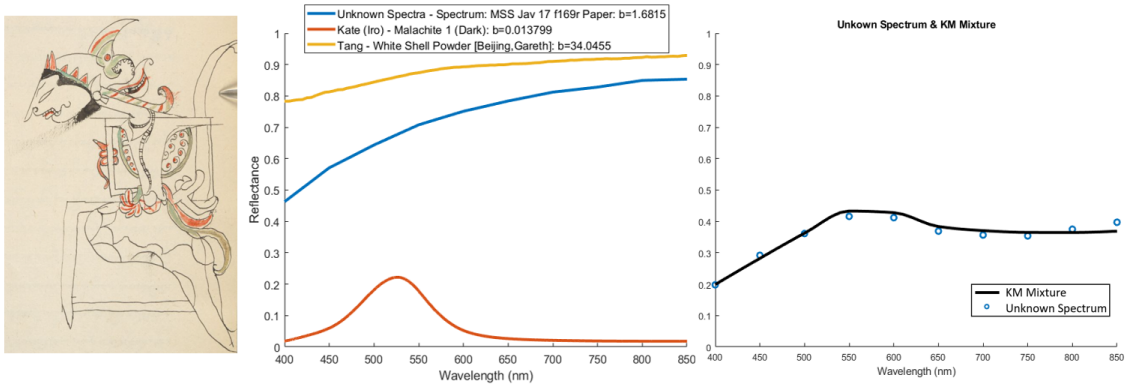

Fig.3.78 Graph showing well-fitting KM mixtures for the green areas on MSS Jav 17 f169r, when a mixture of white, malachite, and paper are used as components.

For the mixture found with the copper green used in MSS Jav 68, it was also found that for the green ship on folio f12r with which the XRF measurement was taken, there was also a strong presence of arsenic and sulphur, suggesting orpiment. Under closer inspection of this painting in the VNIR, it was found that a KM mixture using malachite with an indigo and orpiment mixture fits well with the unknown spectrum for this manuscript, as shown in Fig 3.80, suggesting the use of copper green and the indigo and orpiment mixture commonly encountered in Islamic manuscripts.

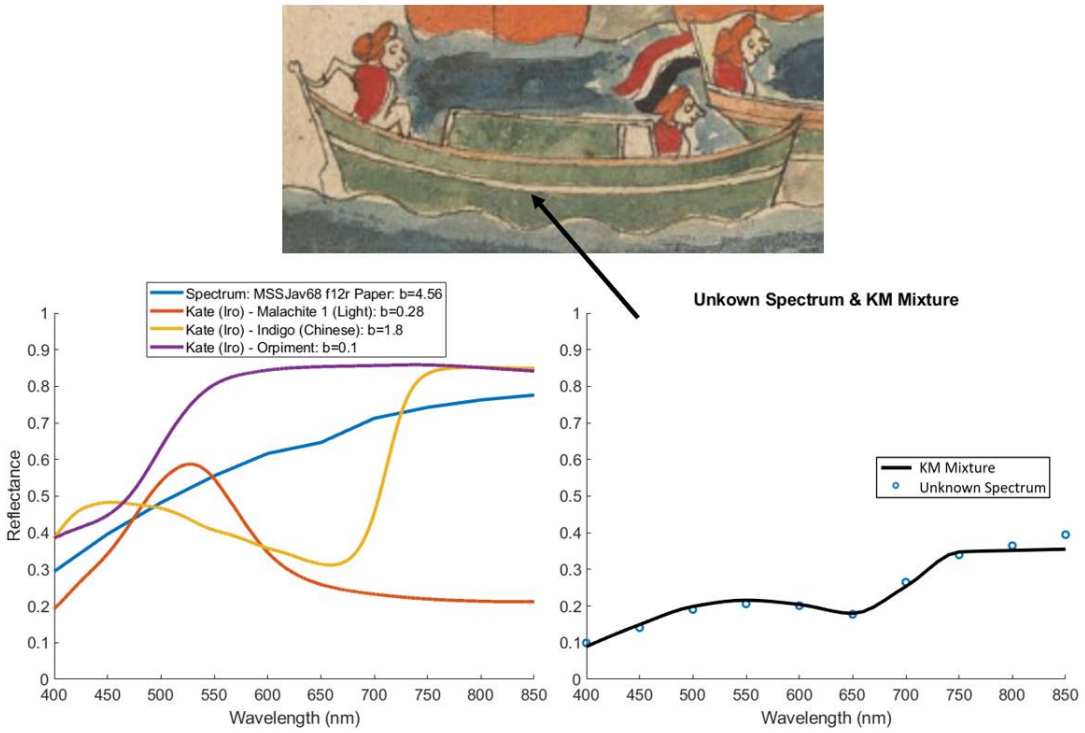

Fig.3.79 Graphs showing how a KM mixture using malachite, indigo, and orpiment together make a well-fitting spectrum to the green used in MSS Jav 68 f12r

Malay Manuscripts

Within the Malay part of the collection copper green pigments were found only two times, where in manuscripts MSS Malay D4 and Or 15227, evidence of could be found through both XRF and VNIR analysis, though the greens used are likely mixed with other pigments, as demonstrated in Fig 3.81.

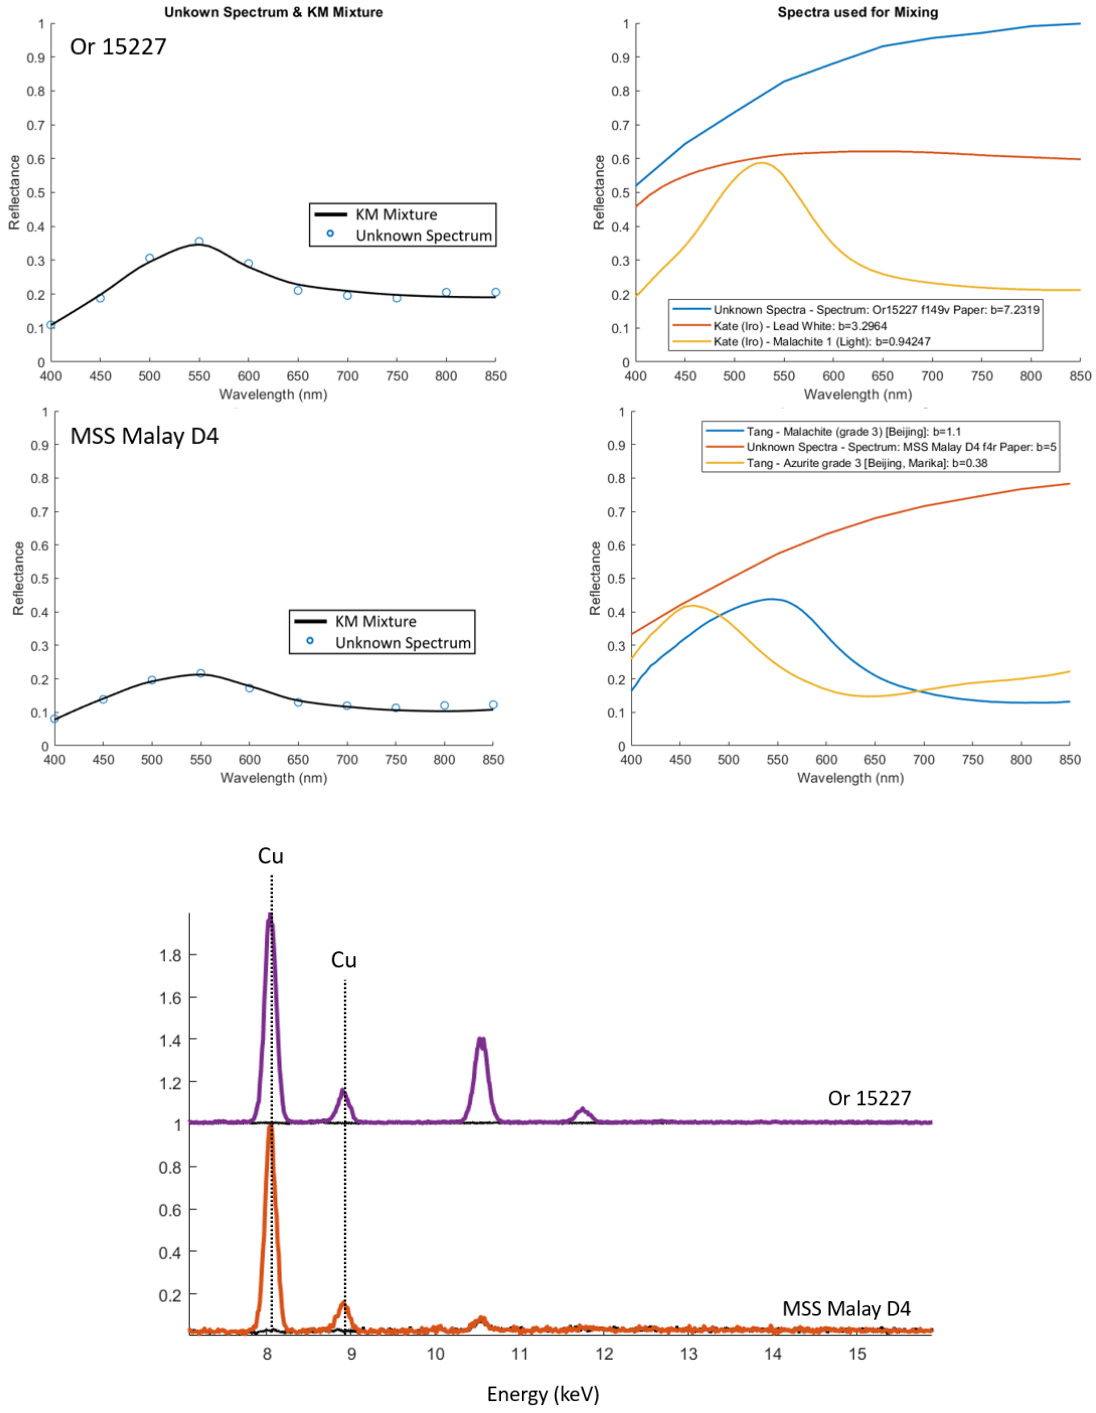

Fig.3.80 (Top) KM mixtures showing well-fitting mixtures when using malachite and paper with other components. Additional lead white may be used in Or 15227, and impurities of azurite in MSS Malay D4. (Bottom) XRF spectra corresponding to both manuscripts are also shown, indicating the use of Copper (Cu).

Sulawesi Manuscripts

Copper-based green pigments are also used in all three Bugis manuscripts within the collection, with each possessing spectra indicative of malachite in both the XRF and the VNIR. Extracting XRF spectra from Or 8154, Add MS 12363, and Add MS 12372 showed the presence of copper in different green regions on all three manuscripts, as demonstrated in Fig 3.82.

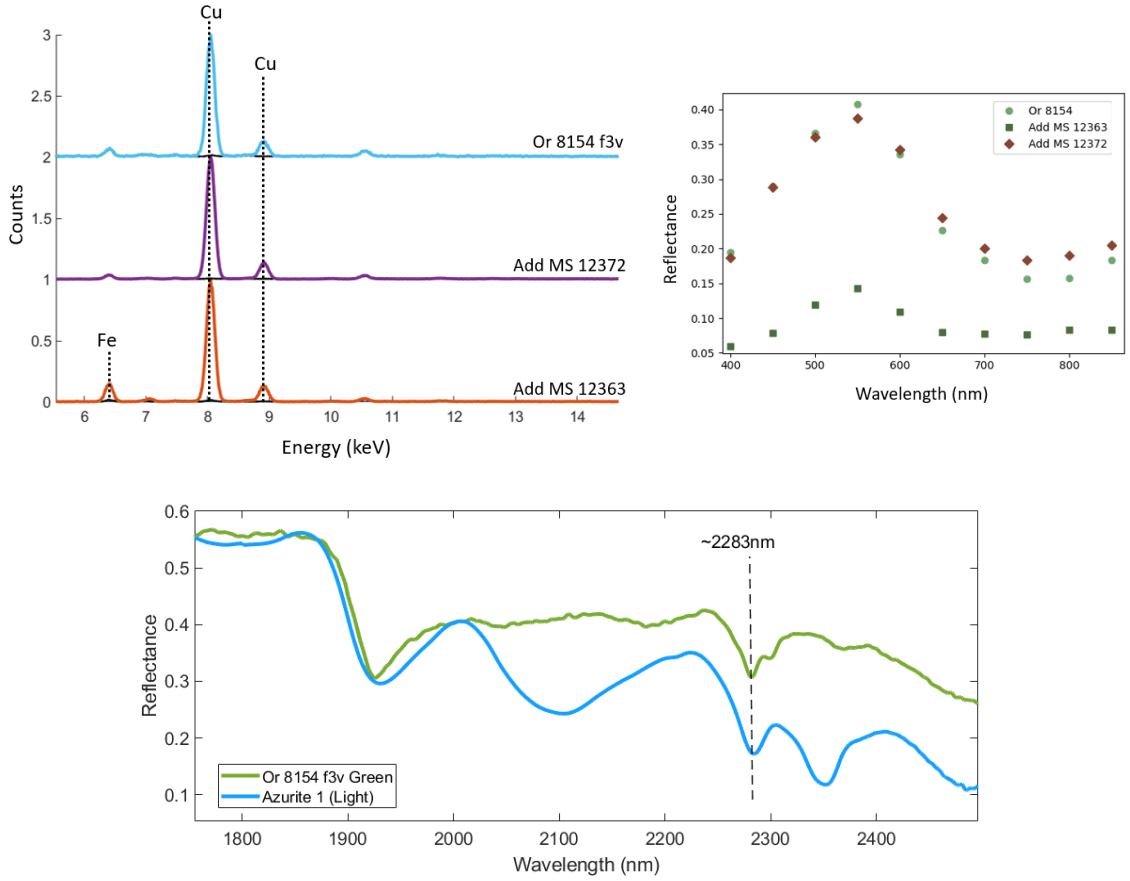

Fig.3.81 (Top Left) XRF spectrum indicating the use of copper in multiple green areas within Bugis manuscripts. (Top Right) VNIR PRISMS spectra indicating the typical malachite features for the three regions. (Bottom) SWIR spectrum showing similar features to Azurite for the copper green in Or 8154 f3v.

When also analysing the VNIR spectra taken from the different manuscripts, clear absorption features can be found for the different manuscripts that correspond well with malachite as a placeholder for copper green, however, the intensity for Add MS 12363 is noticeably lower, owing to the much darker colour of the green areas on the manuscript. This darker green is also copper based, however, and appears to be mixed with another colour to give it a much darker hue. With there being a strong presence of iron (Fe) detected in the XRF signal along with the copper (Cu), and the understanding that the Sulawesi manuscripts all tend to use iron-based inks (see 4.7.1), there is a possibility that the darker colour is achieved by adding iron-gall ink into the malachite, however without further study of the mixture components using additional techniques this cannot be confirmed, as there are also other iron-based pigments which could be used to darken the green. Due to the unique colour of the green found in Or 8154, the green hill was also

investigated further using SWIR FORS (also visible in Fig. 3.82), however neither malachite, atacamite, or Verdigris could fit with the mixture, suggesting an alternative pigment. There was a feature which somewhat corresponded with azurite, at ≈2283nm, however the other features were not present. This could be due to the paper or binder, which may imply the green is a mixture of azurite and a yellow, however without closer study using other techniques this cannot be verified.

3.2.5.2 Indigo + Orpiment

The most common green mixture seen throughout the collection is one produced by mixing indigo with a yellow, which in most cases is most likely to be orpiment. The mixture is easily recognisable due to its green colour, VNIR reflectance spectrum which contains absorption features for both orpiment and indigo and XRF spectrum which contains high levels of arsenic (As) and sulphur (S). However, it can also be confirmed using Raman spectroscopy, where orpiment can normally be detected in the green mixture.

Javanese Manuscripts

The most prominently used green throughout most of the Javanese manuscripts is this indigo orpiment mixture, and it is often found in differing concentrations to achieve different hues. One potentially convenient way of identifying this mixture in the collection was with XRF point measurements. However, as mentioned in the previous section covering yellow colours, it can be difficult to distinguish between different arsenic-based yellow pigments, and this is especially true in mixtures where the VNIR spectrum is altered so there is no clear indication of which yellow is used, therefore later, Raman spectroscopy was also implemented. When analysing the XRF spectra recorded for different Javanese manuscripts where this mixture was suspected to be used, both arsenic (As) and sulphur (S) could be detected, as demonstrated in Fig 3.83.

Furthermore, using KM theory to simulate different mixtures for spectra extracted from the same green regions where point XRF measurements were taken, allowed for good fitting of KM model mixtures with different concentrations of indigo and orpiment. Some examples of these spectra demonstrating the greatest variety of spectral shape can be seen in Fig 3.84. While the use of XRF and VNIR can be quite convincing of the mixture containing orpiment as the yellow, more definitive confirmation can be made that the presence of As and S points to the existence of orpiment by using Raman spectroscopy. Where features indicative of orpiment could be detected in multiple manuscripts, as shown in Fig 3.85.

As Java had many trade connections with countries which could source both indigo and orpiment, it comes as no surprise that many of the Javanese manuscripts used a mixture of the two pigments to achieve different green hues and shades.

Malay Manuscripts

With regards to the manuscripts in the Malay part of the collection, the indigo-orpiment mixture was found only in two separate manuscripts. In Add MS 12379, the mixture can be found in the opening double frames on pages f1v and f2r, and in MSS Malay B3 it exists in the frames on f36v and 37r. Extracted PRISMS spectra from the two regions clearly show the features indicative of the mixture as seen previously. With XRF being used to confirm that the yellow used is arsenic-based and therefore likely orpiment, due to the presence of As and S. KM theory was used to produce a simulated mixture which results in a good fitting spectrum to the extracted PRISMS reflectance data for greens in both manuscripts. The XRF and KM-fitted spectra can be seen in Fig 3.86.

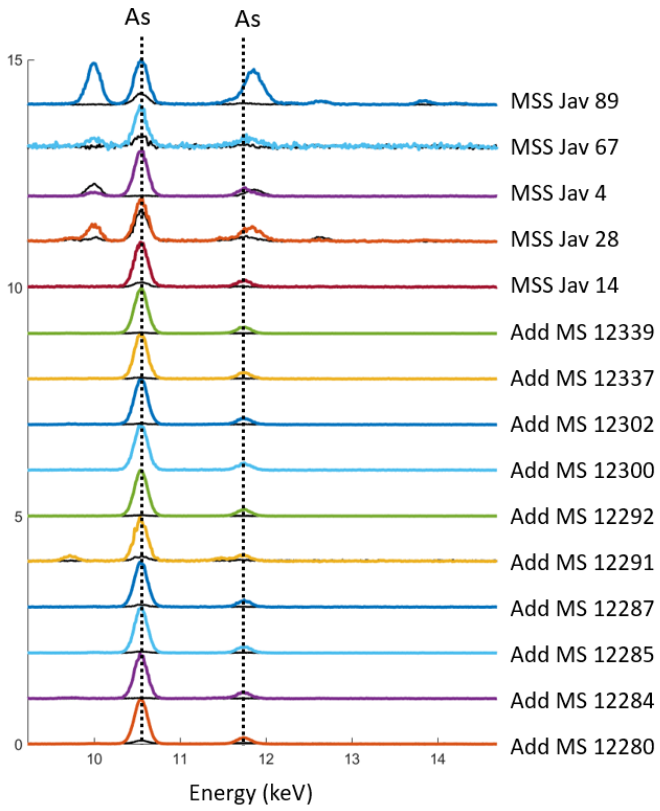

Fig.3.82 XRF spectra taken from green regions in many different Javanese manuscripts, where the presence of arsenic (As) indicates a potential use of orpiment with indigo

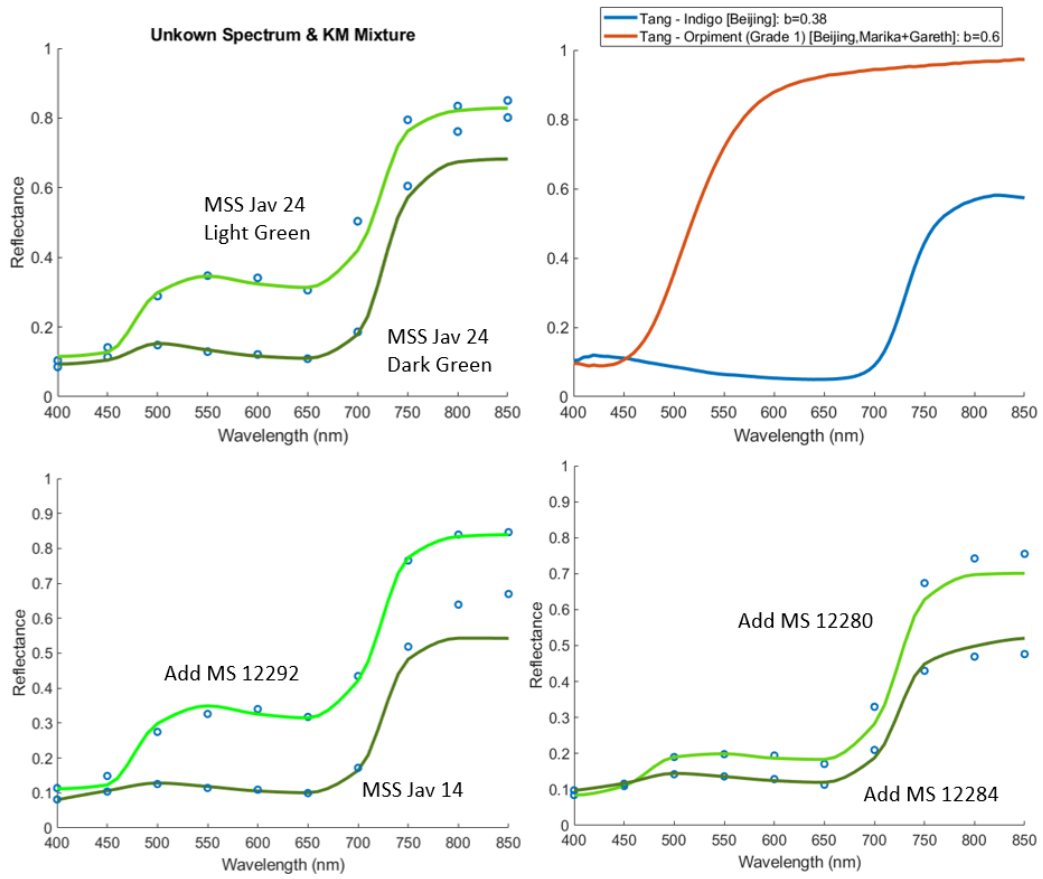

Fig.3.83 KM mixtures for multiple Javanese manuscripts indicating that the use of indigo and orpiment fits well with many of the green extracted spectra from differently coloured green regions.

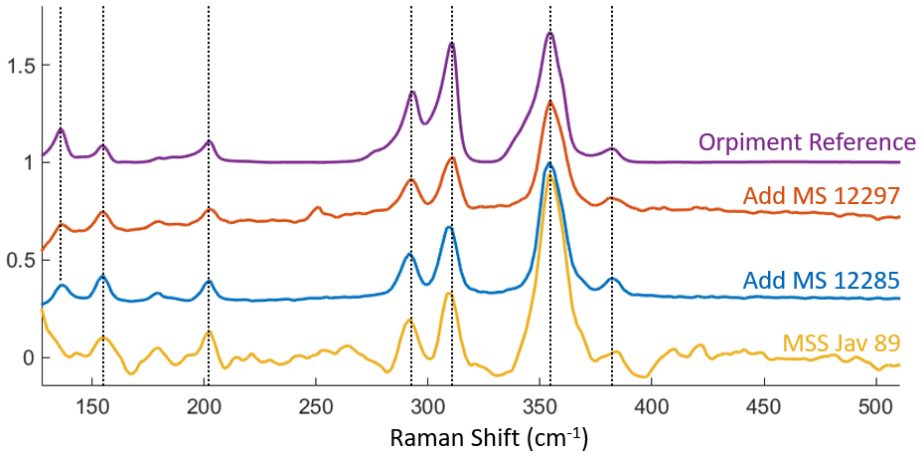

Fig.3.84 Raman spectra for green indigo and orpiment areas seen in multiple Javanese manuscripts. Each share the same features with the orpiment reference taken from (RRUFF, 2023b)

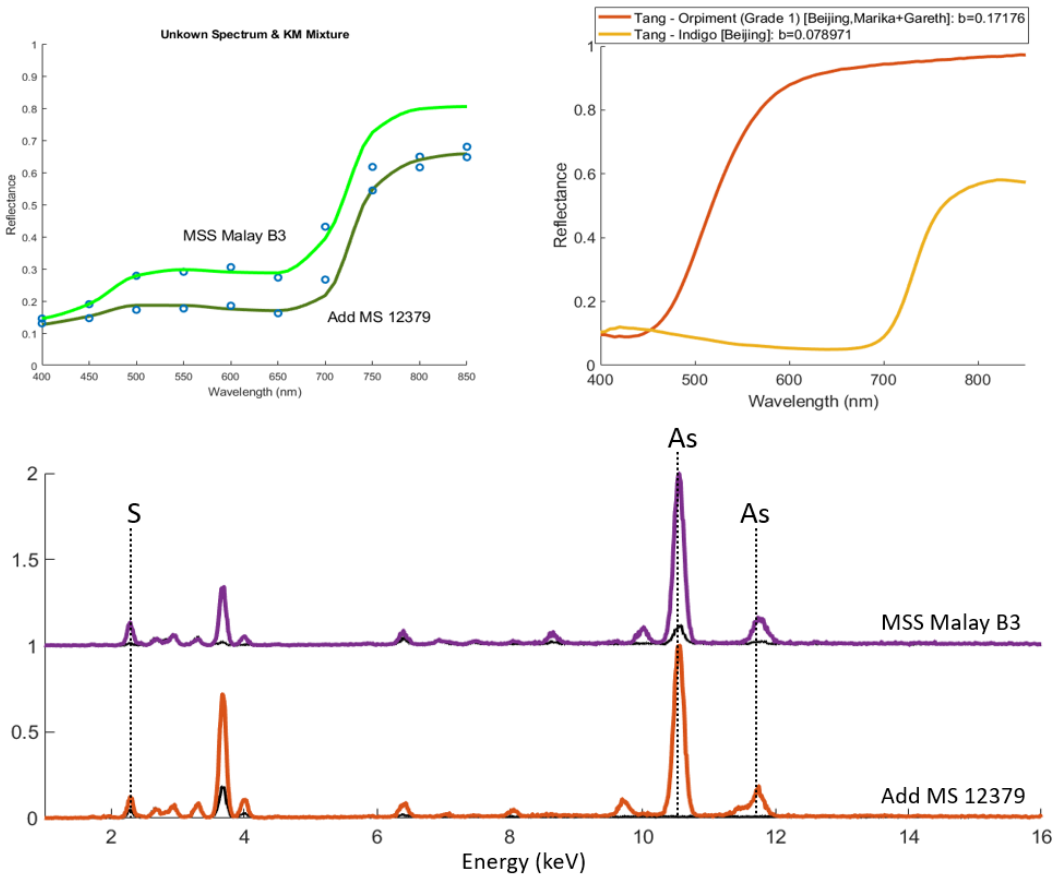

Fig.3.85 (Top) graphs showing that the same indigo and orpiment mixture fits well for Malay manuscripts MSS Malay B3 and Add MS 12379.) Both are also confirmed to likely possess orpiment due to the presence of both arsenic (As) and sulphur (S) in the normalised XRF spectra (Bottom).

Sumatran (Acehnese) Manuscripts

Though most Acehnese manuscripts do not typically use green colours within their illumination, Or 16769 is considered the exception, where a green pigment is used in a set of double frames on pages f6v and f7r. Again, the joint use of XRF and PRISMS reflectance spectra were used to identify the constituent materials used in the green paint, with which the use of indigo and orpiment created a good fitting virtual spectrum using KM modelling, the spectra of which can be seen in Fig 3.87.

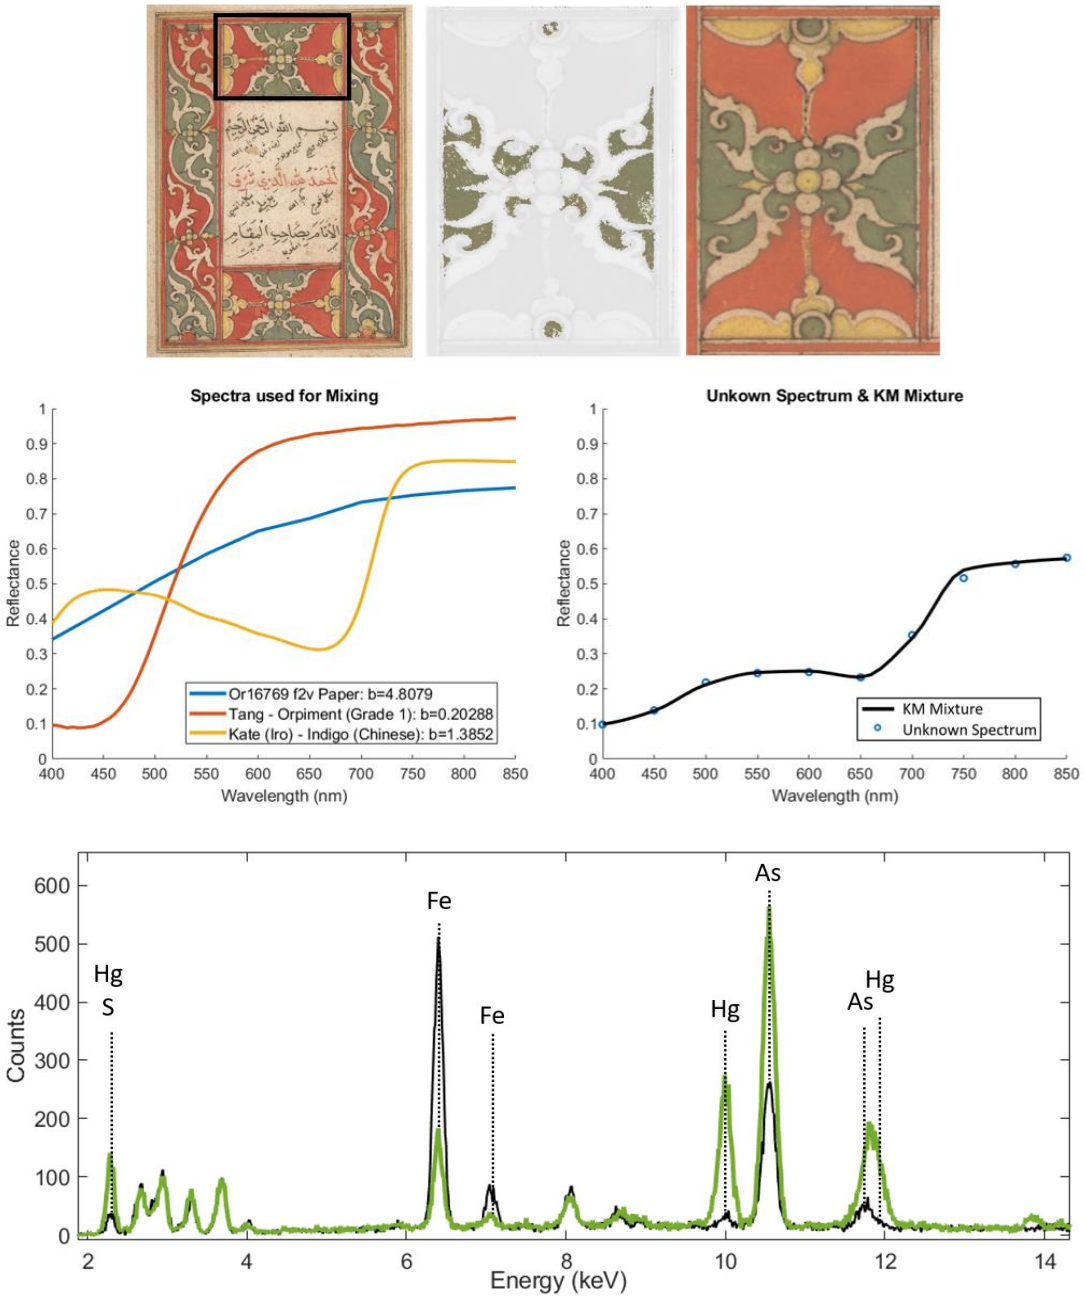

Fig.3.86 KM mixture showing that indigo and orpiment can be used as components with paper to create a well-fitting mixture to the green cluster in the Acehnese manuscript Or 16769, on folio f6v. XRF spectra of the green and paper (black) can be seen showing presence of arsenic (As) indicative of orpiment.

3.2.5.3 Prussian Blue + Organic Yellow (Gamboge) (Hooker's Green)

Javanese Manuscripts – MSS Jav 67

There is evidence of a Prussian blue and gamboge mixture is used on MSS Jav 67, where multiple characters were consistently illuminated with watercolour style paints as opposed to the more opaque paints seen in more traditional manuscripts. Hooker's green tends to be found on multiple pages but was identified using PRISMS, XRF and ER-FTIR on f12r, where a small bird is illuminated with green in its wings. Firstly, a PRISMS reflectance spectrum was used to rule out some of the common green pigments used throughout the collection, such as the indigo-orpiment mix seen in many Javanese manuscripts. Under closer inspection, the spectrum appeared to follow a shape commonly seen for copper-based green pigments, such as malachite, however, the increase in reflectance towards the red and infrared parts of the spectrum suggested otherwise. A point measurement taken from the same location using XRF showed that there was no copper, but there was an increase in the counts for iron (Fe) and potassium (K), suggesting the use of Prussian Blue with a yellow pigment. With no other significant XRF peaks, it was determined that the yellow was most likely organic. Using KM modelling to simulate the mixtures for several different yellow pigments with Prussian blue resulted in multiple spectra which could fit the unknown green mix, therefore ER-FTIR was also used to perform a final identification, where evidence in favour of gamboge being used could be found, due to the features at approximately 1597-1599cm^{-1} and 1730-1750cm^{-1} which match closely with those seen in (Biron et al., 2020) and demonstrated in Fig 3.88.

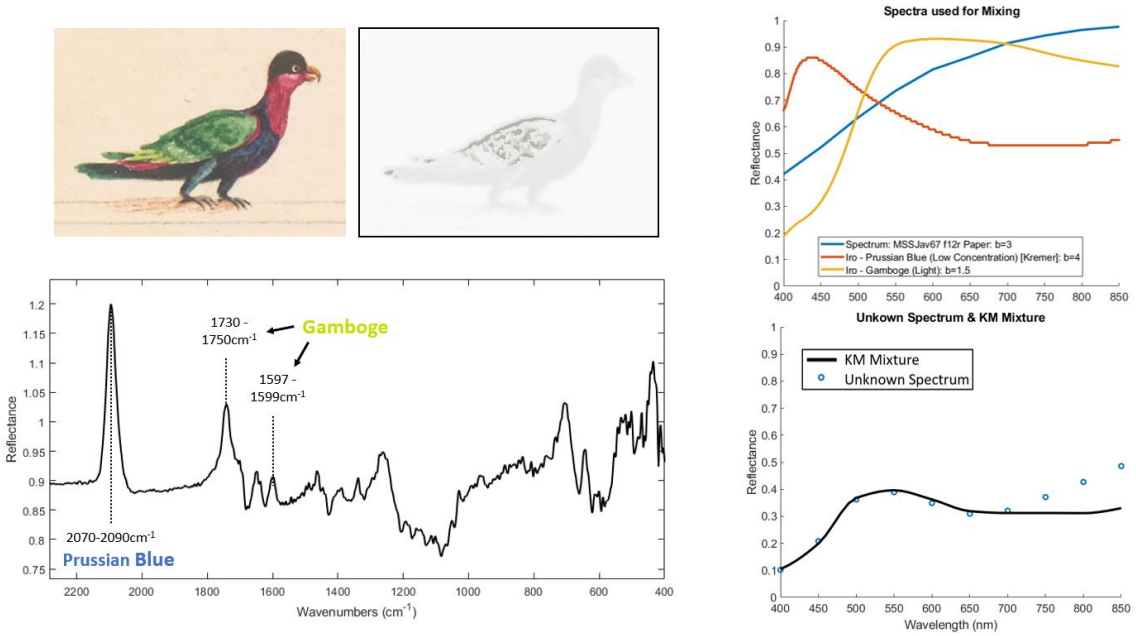

Fig.3.87 Multiple graphs showing the (Top Left) RGB and true colour cluster map for the green region of the bird's wing. (Right) KM mixture indicating a mixture of Prussian blue and gamboge fits well to the unknown spectrum. (Bottom Left) ER-FTIR spectra confirming the presence of Prussian blue, and an organic yellow most likely gamboge.

3.2.5.4 Indigo + Organic Yellow

Javanese Manuscripts – Or 15932

Evidence suggests that the use of a mixture of both indigo and an organic yellow can be seen within one of the Javanese manuscripts in the collection. Or 15932 is another of the Javanese manuscripts which have a more western artistic style and uses watercolour type paints instead of traditional colours which are opaquer. It is expected that some of the pigments used may differ from the earlier 18th-century manuscripts due to their different style, however, the mixture of indigo and organic yellow appears to be unique to this manuscript, and so is explored only within this example. The PRISMS data extracted from the outer green leaves of a painted tree on f21v possessed a reflectance spectrum which possessed features similar to the very common indigo-orpiment mixture seen in many Javanese manuscripts, as seen for the cluster map and KM fitted spectra shown in Fig 3.89.

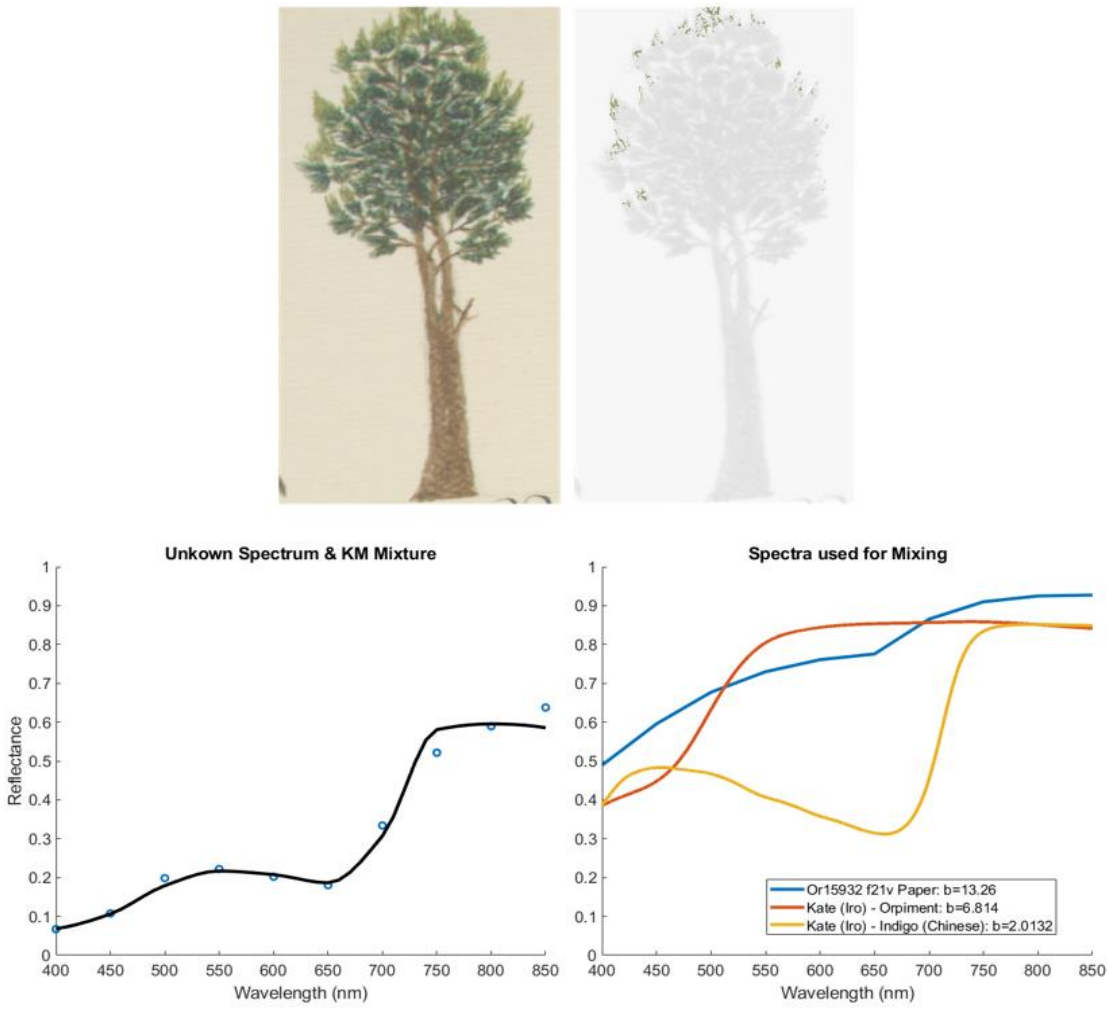

Fig.3.88 KM mixture showing that a combination of indigo and an orpiment can be used to create a good fit to the unknown green found in Or 15932 f21v. However, this cannot be the correct mixture as XRF shows no arsenic.

However, the XRF spectrum for the same region showed no presence of orpiment, with the only changes from paper being an increase in counts for calcium (Ca), suggesting that the yellow part of the mixture must be an inorganic pigment, such as gamboge seen for the previous green Prussian blue and gamboge mixture. ER-FTIR data was taken for this green mixture to investigate further, however no clear matching features with any references for different potential organic yellows such as gamboge, saffron, turmeric, or Indian yellow could be found, so further study would be required to be sure of which material is used.

3.2.6 Brown Pigments & Mixtures

3.2.6.1 Indigo + Vermillion + Orpiment

Javanese Manuscripts

In some of the Javanese manuscripts, brown paints are implemented in the decorative frames often seen at the opening folios, or in the illustrations of different characters. Under close examination of the PRISMS reflectance spectra, it is difficult to identify the brown with any natural pigment, however, there does appear to be a spectral feature which may suggest the presence of indigo. This however does not provide any insight into which other pigments may have been used to create the brown mixture, but with the use of XRF, constituent inorganic elements can be determined.

Point measurements taken from the brown decorative pillars on Add MS 12280 were used to perform analysis for all the manuscripts which had this mixture, or spectra very similar to it. XRF showed evidence of mercury (Hg), arsenic (As), and sulphur (S), in addition to the materials present on the paper, suggesting the use of vermillion, and orpiment with the indigo, which fits well as blue, yellow, and red can make brown shades (see Fig 3.90).

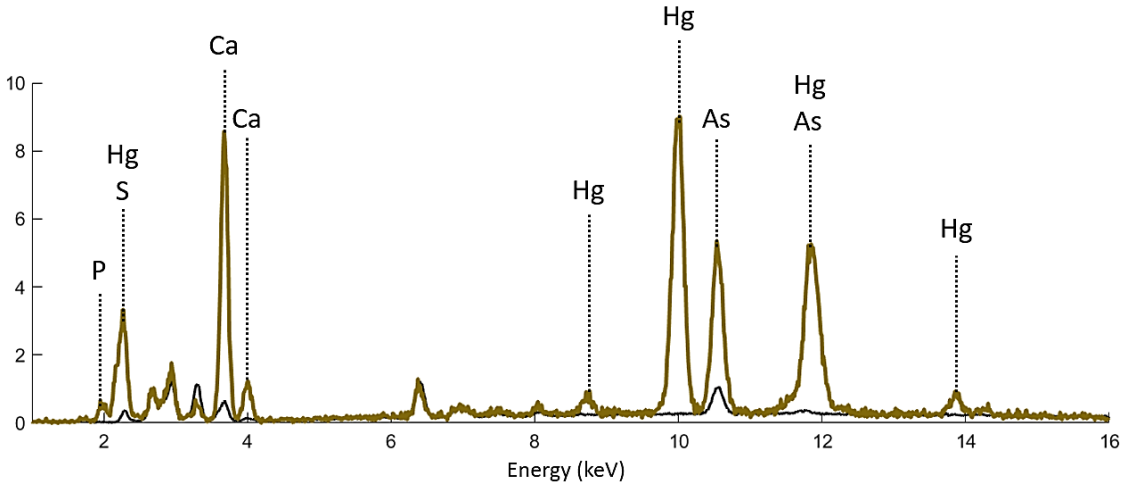

Fig.3.89 XRF spectrum showing strong peaks for mercury (Hg), arsenic (As), Calcium (Ca), Sulphur (S) and phosphorus (P), indicating the use of orpiment, vermillion, and bone ash white to make a brown.

If the assumption is made that the presence of Hg, As, and S is a result of mixing indigo with vermillion and an indigo/orpiment green, a Kubelka-Munk fit can be made and compared against the PRISMS spectrum to investigate the likelihood that these materials are the components of the brown paint. Such a fit can be seen in 3.91, where paper and a white pigment are added into the mixture, to account for any transparency toward the NIR and also mimic the detected presence of bone ash white detected in the XRF.

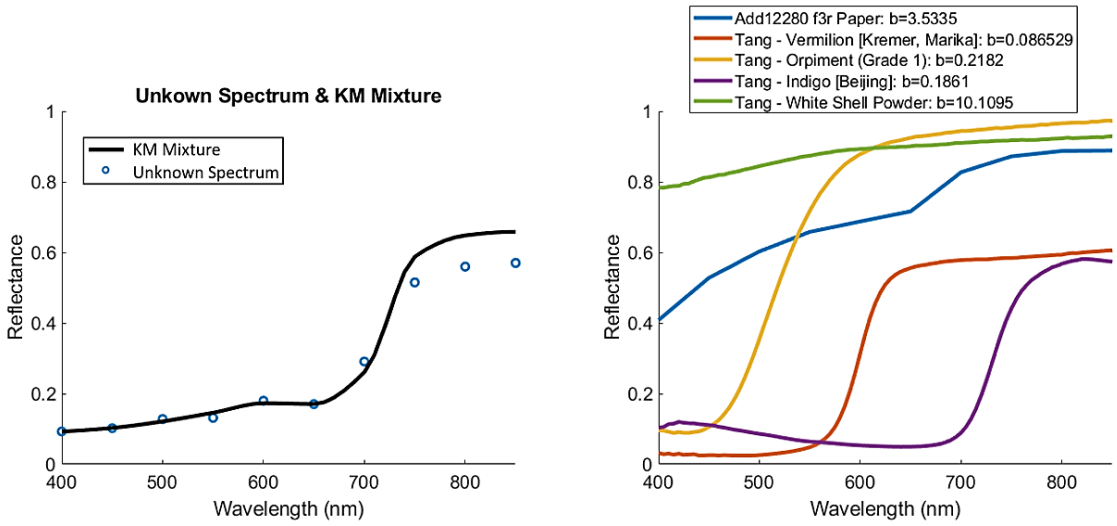

Fig.3.90 KM mixture showing that a combination of paper, vermillion, orpiment, indigo, and white are used to make a brown in Add MS 12280 f3r

The results of the KM-modelled mixture demonstrate that a combination of these three materials does have a good fit, suggesting that this is the mixture most likely used to achieve the brown mixed paint seen in Add MS 12280, and therefore by extension other manuscripts which may share the same VNIR and XRF spectrum.

3.2.6.2 Prussian Blue + Vermillion + Orpiment

In the Malay manuscript Or 13295 there is evidence to suggest that a dark red/brown shade on the double frames within folios 190v and 191r is created by mixing Prussian blue, Vermillion, and orpiment. The XRF spectrum presents features typical of pigments seen over the rest of the manuscript, with strong signal present for pigments such as vermillion (Hg, S), Orpiment (As), and Prussian Blue (Fe), and when using the three pigments together to create a virtual mixture using KM mixture modelling, the three create a spectrum which matches almost perfectly with cluster 6741, which represents a significant number of the brown pixels in this manuscript. The associated KM mixture spectra and XRF spectrum can be seen in Fig 3.92. While this is a good fitting mixture however, other point measurements, such as ER-FTIR, will likely be needed to confirm it in future.

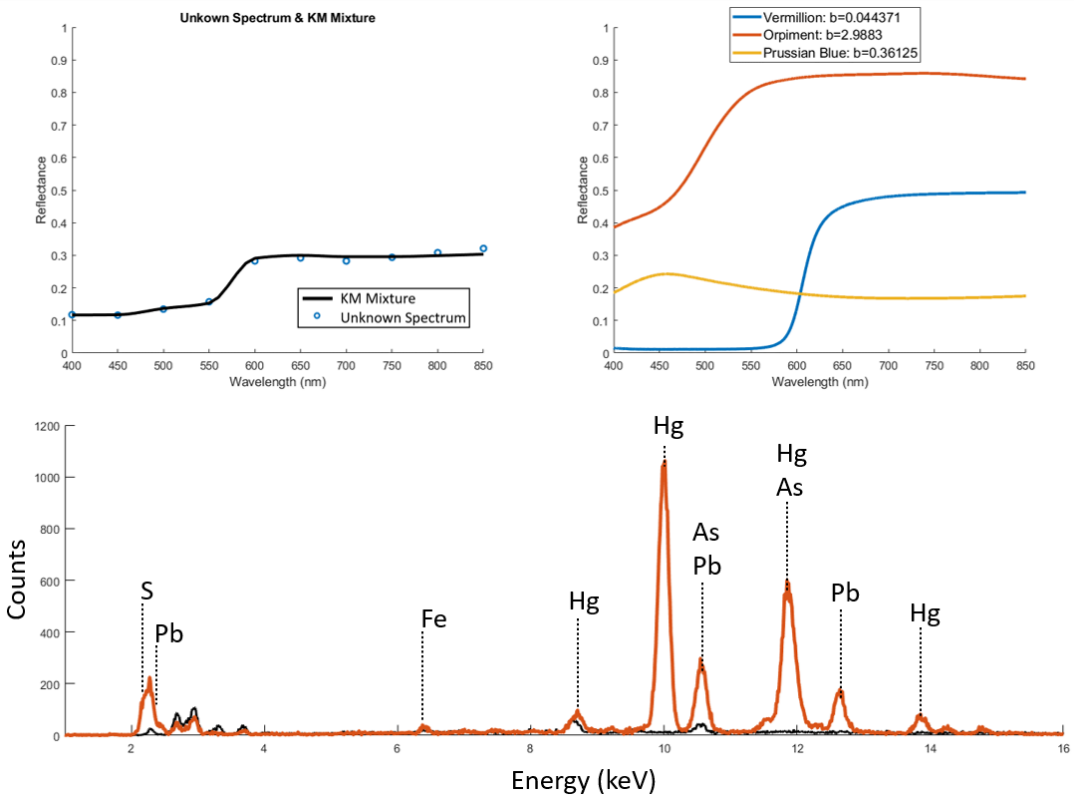

Fig.3.91 KM mixture showing how orpiment, vermillion and Prussian blue can be used together to create a well-fitting spectrum that matches with the dark reddish-brown frame of Or 13295 f190v. Support by XRF (bottom) which shows elements indicative of all three pigments: Mercury (Hg), arsenic (As), sulphur (S) and iron (Fe). Notice the addition of lead (Pb), which is likely indicative of lead white used to lighten the mixture.

3.2.6.3 Vermillion + Yellow Ochre + Brown Ochre

There is evidence to suggest that a mixture of vermillion, yellow ochre and brown ochre is used on Or 15932 f72v for the tree which dominates a significant portion of the page. However, when monitoring the XRF spectrum for the brown area, there was a strong presence of mercury (Hg) and sulphur (S) with no clear change in iron (Fe), suggesting that the only inorganic pigment present in the mixture is vermillion. As iron (Fe) would have been detected in large amounts if ochre had been used it was first concluded that the brown mixture may be an entirely different mixture, however when analysing FORS spectra extracted from a brown location, the VNIR spectrum matched more closely with a mixture of vermillion with red, yellow and brown ochre than anything else. The XRF spectrum and associated KM fitting spectrum used with FORS data in the range 400-1000nm can be seen demonstrated in Fig 3.93. This is likely a result of the XRF and FORS not being recorded from quite the same location, resulting in one of the small redder appearing spots having a greater effect on the XRF signal than for the FORS spectra. Either way, the KM fitting performs better when implementing ochre into the mixture, so it is likely that in certain parts of this brown tree, ochre is used.

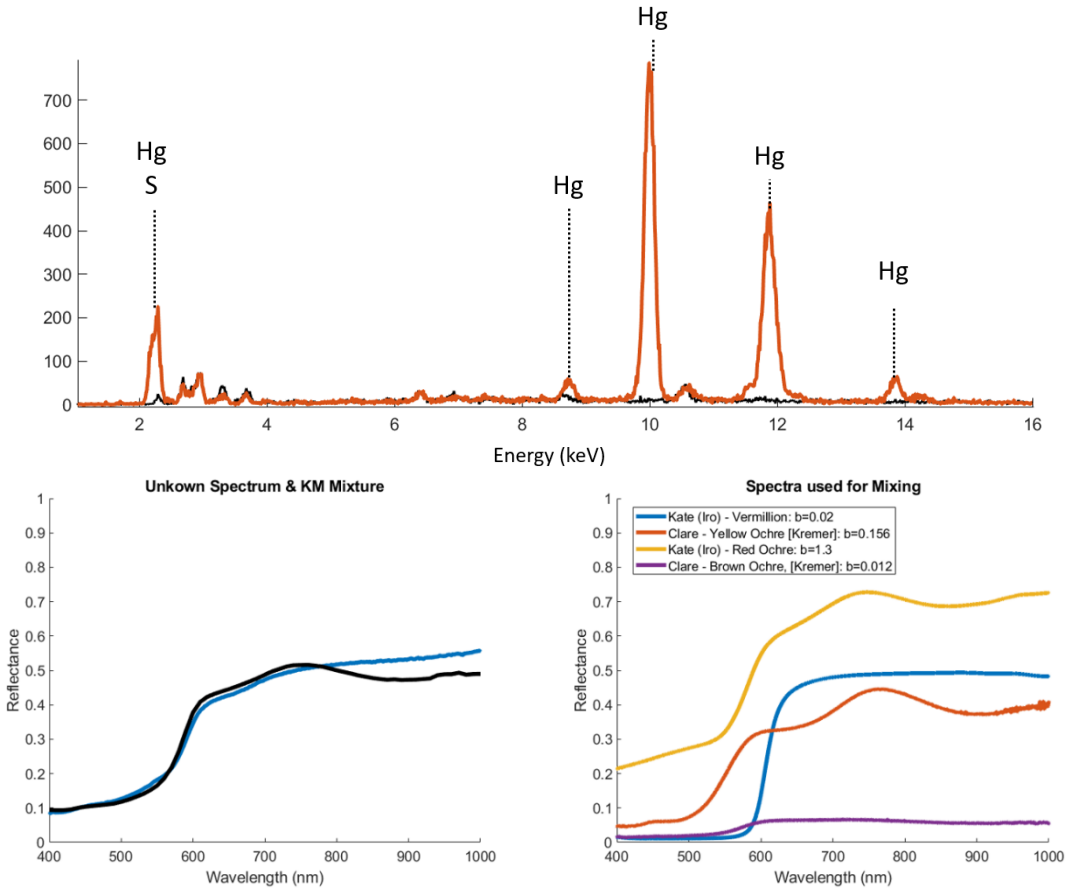

Fig.3.92 (Top) XRF spectra indicating the presence of vermillion, via the clear peaks of mercury (Hg) and sulphur (S), though for a more brown spot, taken using FORS, spectra closer to that of a mixture of vermillion with red/brown ochre could be detected, as is indicated by the KM mixture.

3.2.7 Black Pigments & Inks

As the VNIR reflectance spectra are not particularly useful for the analysis of most black colours and inks, XRF and NIR spectral image bands were primarily used to aid in their analysis and identification. Carbon black is noted to have been used in some of the manuscripts seen in this collection, however as Raman spectroscopy was not used to confirm this, and other detection methods used on inks within this study cannot detect carbon, the study was primarily focussed on other inks and black pigments.

3.2.7.1 Iron-Based Ink (Iron Gall)

Within the collection, there is evidence of different iron-based inks being used, with the most likely being iron gall ink, which has been used throughout many parts of the world for centuries and can be found in many manuscripts with different cultural origins. Iron gall was the primary pigment used within Europe from as early as the 5th century, was used regularly in the Islamic world and can still be found in use today, making it very possible it to have appeared in 18th-19th century illuminated manuscripts within maritime Southeast Asia, especially considering the influence Europe and Islam nations would have had at the time. Furthermore, by the 18th century, mass production of iron gall ink was more commonplace due to the rise of industrialisation in the west, increasing the chances of the ink being readily available among maritime trade routes. Iron gall ink production would often involve the use of iron (II) and gallic acid (Raggetti, 2021) , (normally extracted from plant materials such as gall nuts, which is where the ink owes its name), though there are many variations and different recipes for producing it, which can vary widely based on the cultures or times throughout history with which the ink was produced (Lemay, 2013). It is therefore expected that there may be some variations in the use of iron gall ink in maritime Southeast Asia due to the many different connections that the region had with other locations during the 17th-19th centuries. If iron gall ink can be detected in different manuscripts, it would be interesting to understand how their provenance may link with its use and therefore how different cultures may have had an impact on the production of illuminated manuscripts in different regions. Iron gall ink can be detected in many ways, however, due to the main ingredient of the ink being iron, it can be simple to find evidence for its existence in different manuscripts by using XRF spectroscopy, which would provide energy peaks at approximately 6.4keV and 7.06keV, though other techniques such as FTIR and Raman can be used (Raggetti, 2021). Furthermore, iron gall can also be detected using PRISMS as it becomes transparent towards the NIR parts of the electromagnetic spectrum (Havermans et al., 2003). Due to this simplicity, a combination of XRF and spectral imaging was used to identify iron gall over the collection. In total, iron gall was detected in 10 manuscripts, with the potential for it to also exist in a few more, though these were unclear. The manuscripts which were seen to potentially contain iron gall ink are as follows:

- Add MS 12292 (Inner Writing)
- Add MS 12297 (Inner Writing)
- Add MS 12300 (Inner Writing)
- Add MS 12363 (f63v - Smudges)
- Add MS 12372 (f66r Circle)
- F-148-4-f106 (Smudges)
- Mackenzie Pv. 42 (Writing)
- MSS Jav 67 (Writing)
- MSS EUR E378-1 (Writing)
- Or 15026 (Writing)

In Or 9333, there is also some evidence suggesting iron gall ink had been used previously but has since been removed. This can be seen in folio f11r, where in the bottom corner a removed piece of writing can be shown to become more transparent at 850nm, as demonstrated in Fig 3.94. XRF also showed significant content of iron in the Sumatran manuscripts (see Fig 3.95), however, the 850nm bands for these manuscripts did not show any transparency.

Fig.3.93 PRISMS spectral images demonstrating potential iron gall ink regions becoming transparent in the NIR (850nm) on Or 9333 f11r.

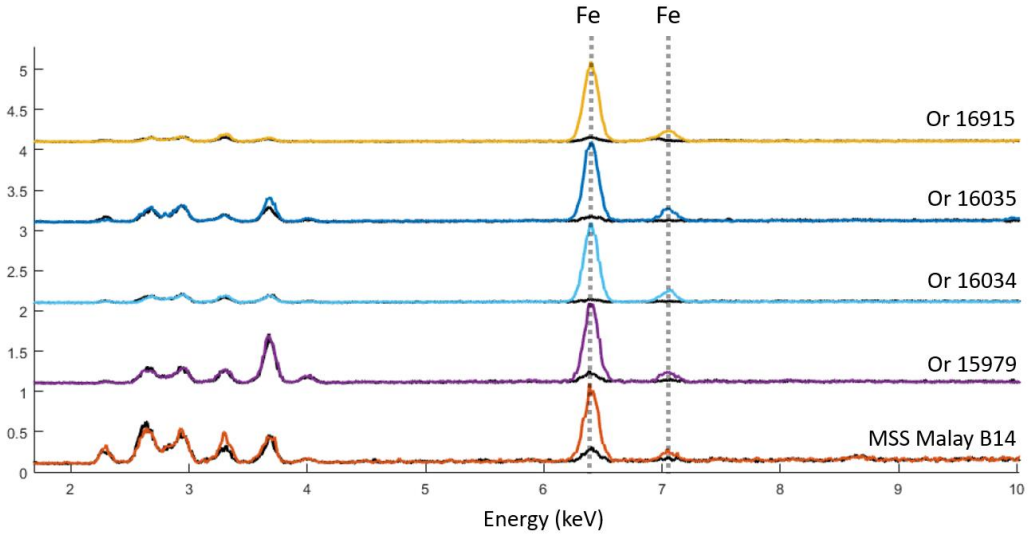

Fig.3.94 XRF spectra showing strong iron presence on multiple Sumatran manuscripts, indicating iron gall ink even though the inks did not become transparent toward the NIR.

3.2.7.2 Copper-Based Inks

In numerous manuscripts, copper can often be detected by XRF in the inks or black-painted areas. This copper presence could be attributed to the use of black copper (II) oxide or copper sulphate (Fichera et al., 2018) as an ingredient in creating ink, however, there is very little literature to suggest this was ever used in the 18th-19th centuries in maritime Southeast Asia. Within the literature, there has been considerable study into the existence of copper in other inks such as impurities in iron-gall, or its potential use within renaissance printing inks to enhance colour (Stanley, 2018). However, in most manuscripts with signs of a copper-based ink seen within this collection, the copper presence in the XRF spectrum tends to be higher in XRF counts than any other element and often is accompanied with little to no iron.

The following manuscripts were shown to have contained significant copper presence within their inks, where those with additional detected elements are mentioned in the right-hand side column:

- Add MS 12302
- Add MS 12312
- Add MS 12337
- Add MS 12339
- MSS Malay B12
- MSS Malay B3
- Or 8154
- MSS Jav 17 (+Fe, Pb)
- MSS Jav 24 (+Fe, Ni, Zn, As)
- MSS Jav 28 (+Fe, As, Zn)
- MSS Malay D4 (+ Trace Pb)
- Add MS 12379 (+Hg, Trace Zn)
- MSS Jav 14 (+As, Trace Zn)
- MSS Jav 16 (+Zn)

For the manuscripts which possess iron content (MSS Jav 17, MSS Jav 24, and MSS Jav 28) in addition to copper, the source of copper in the XRF could be due to its use during the creation of an iron gall ink which is now mixed with another ink, though in this circumstance it would be strange for these manuscripts to have an iron content secondary to the copper. However, for the other inks where there is no detectable iron content, and copper still exists as the primary element, the question of which materials may be used still needs to be answered. It is more likely for these inks that the presence of copper is due to there being contaminants within a carbon black, which is likely to be used on manuscripts due to its well-known use within Asia (Singh, 2020; Stanley, 2017) and the Islamic world (Kantoğlu et al., 2018; Knipe et al., 2018). The copper could potentially also exist as a result of the direct use of copper sulphate or black copper oxide as ink, however further study would have to be performed to determine this as there is little evidence to suggest that this was a normal practice within the region.

3.2.7.3 Bone Black – Add MS 12291

Within one manuscript evidence was also found to suggest the use of bone black, due to the presence of calcium (Ca) peaks which are much higher in XRF counts than what can be detected in the paper, as shown in Fig 3.96. Potassium (K), Sulphur (S) and silicon (Si) presence within the XRF spectrum is likely due to impurities in the mixture, but the presence of calcium and potentially trace phosphorus (P), clearly links the ink used in this manuscript with bone black which typically has the elemental composition $Ca_3(PO_4)_2 + CaCO_3 + C$. This however would need further confirmation using alternative techniques to be sure.

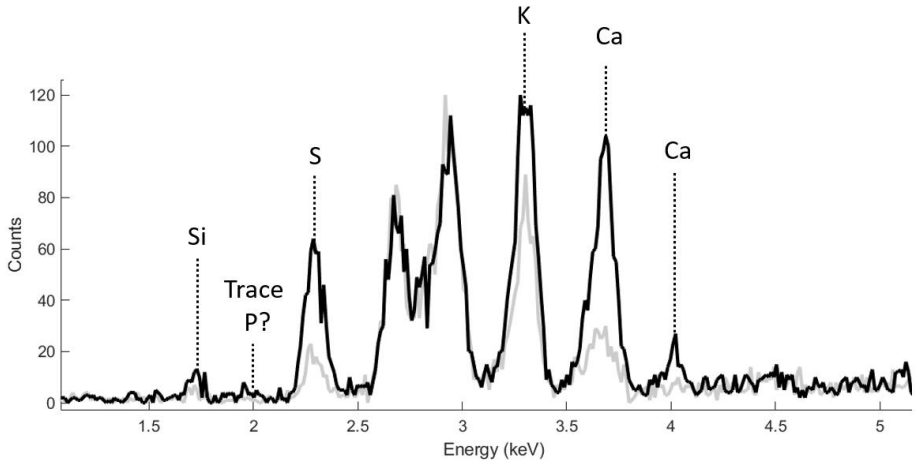

Fig.3.95 XRF spectrum indicating the use of bone black in Add MS 12291, indicated by the presence of calcium (Ca) and perhaps some trace phosphorus (P).

3.2.7.4 Mercury-Based Black – MSS Jav 4

In the Javanese manuscript MSS Jav 4, there are black areas around the outsides of the opening double frames that, when analysed using XRF, showed a strong presence of mercury (Hg) and sulphur (S), which highly suggests the use of vermillion (see Fig 3.97). While vermillion is known to blacken over time (Spring et al., 2002), this is an unlikely cause for the black areas on this manuscript, as even though the pages have seen some degradation over the years, other parts of the page still possess the same bright vermillion red which would have been intentionally applied onto the paper. This may imply that a use of black vermillion in this manuscript is intentional.

An alternative, more likely, cause of the presence of vermillion is that the red was painted over a black already present on the manuscript page, instead of being a black pigment itself. There is limited evidence to suggest this, however, visual observation of the manuscript shows that there are certain areas of this black part of the frame that seemingly have red or brownish hues attributed to paint which may be layered over the top. Such an example can be seen in Fig 3.98, which also shows that a combination of vermillion and black (in this case ivory black) could be used to create a KM mixture with paper that fits well to a cluster of these reddish-brown areas.

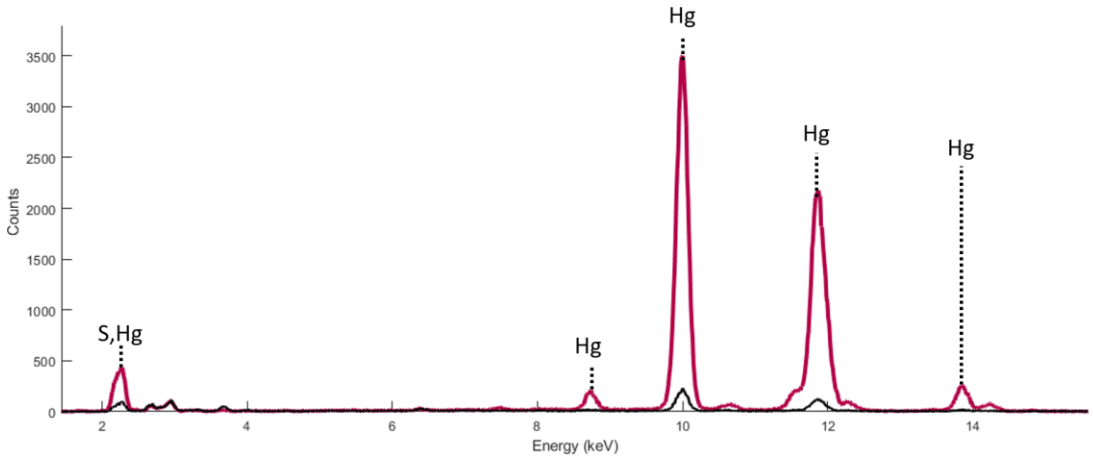

Fig.3.96 XRF spectrum showing strong peaks for mercury (Hg) and sulphur (S) suggesting a use of vermillion in the black parts of MSS Jav 4's opening illuminated double frames.

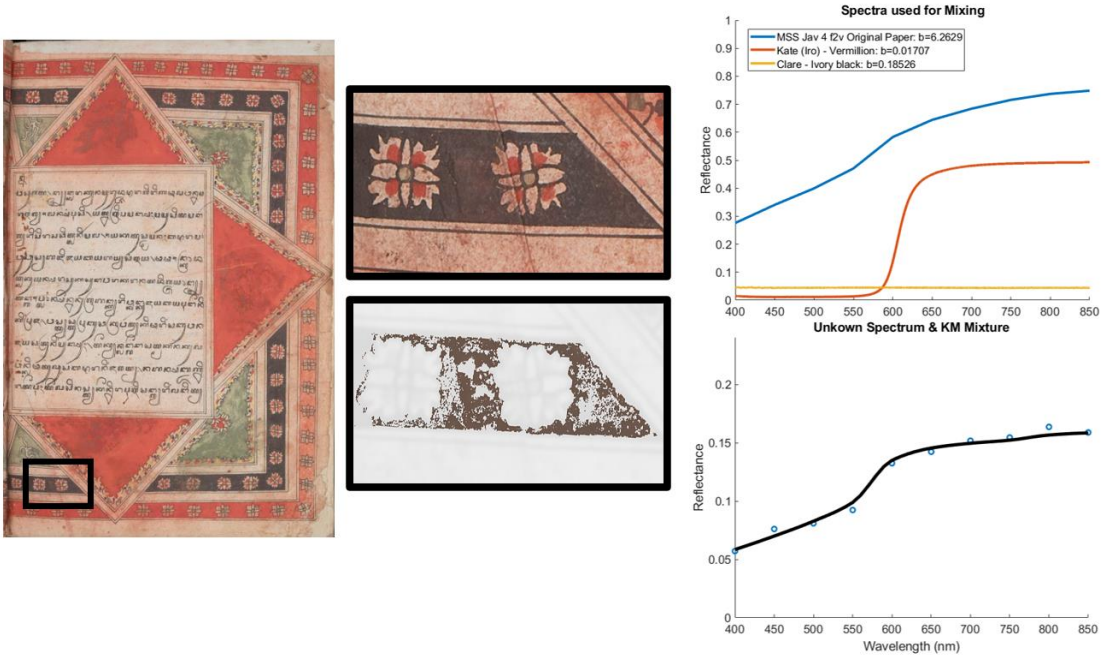

Fig.3.97 KM mixture for a mean cluster spectrum representing the dark vermillion parts of MSS Jav 4, which shows that a mixture of vermillion with paper and black fits well. Possibly suggesting a layering of vermilion over black instead of blackened vermillion.

3.3 Pigment Mixture Mapping

Throughout the previous section many multimodal point measurements, including VNIR spectra extracted from PRISMS data and clustering, were used to successfully identify many different specific pigments and mixtures in varying concentrations over different manuscripts. This provided plenty of information which is valuable to understanding how the use, trade and exchange of artistic materials may have evolved between regions and across trade routes within maritime Southeast Asia. To reiterate, it is essential for this study that large amounts of data are collected as the main purpose of analysing an entire collection is to ensure that a significant enough sample size is used to make statistically reliable conclusions about the pigments and artistic styles used in different regions. However, determining these conclusions is not the only motivation for performing these point measurements, another important reason is also because building reliable connections between colour/VNIR spectral data and the pigments with which these spectra represent can allow for highly reliable pigment mixture maps to be created. Throughout this chapter, cluster grouping was performed simultaneously to pigment identification using the hybrid Kubelka-Munk model informed approach outlined in chapter 2.6.4, resulting in accurate illustrations of pigment mixture distribution over the entire collection.

To therefore perform a final collection-wide summary, different groups created for specific pigment mixtures could be further updated or regrouped and then analysed to determine the distribution of artistic materials used throughout the entire collection. Essentially, this regrouping approach followed the same example shown in chapter 2, where painted areas would be grouped by selecting regions of interest at positions known to be a specific mixture, however instead of focussing on grouping using a single spectral imaging data cube or manuscript, grouping was performed by selecting sets of manuscripts from the collection that were shown to possess a good variety and volume of areas where specific pigments and mixtures used for the illumination were either known outright or extrapolated from the results of grouping which had already been performed. This meant that, for example, if a set of clusters representing vermillion are grouped on one manuscript, in turn grouping vermillion clusters on a second manuscript which has more spectral variation than the first, then this second manuscript may also have its own vermillion clusters added to the same group as well. This process ensured that more complete cluster maps could be generated by attempting to encompass all the spectral variations which may occur on account of there being changes in constituent pigment concentration, paint layer thickness, or substrate. There was however a risk that the limitations of performing the grouping using VNIR clusters would result in misidentified pigment map areas where the constituent artistic materials used to create colours are different, even though the VNIR reflectance spectrum and colour appear the same. This potential for misidentification was the main reason why so many XRF point measurements were recorded for different manuscripts, as at least in this study, if these problems are encountered it is likely that they will be discovered. Some of these discovered

misidentifications are discussed further in the following section (3.3.1), however, it was found that for most materials, pigments and pigment mixtures, the clustering and grouping process performed well.

3.3.1 Clustering & Grouping Limitations/Misidentifications

While the pigment mixture mapping worked well to represent most artistic materials, there were some mixtures which were expected to have incorrect representations in different manuscripts from clustering and grouping. This is unfortunately a consequence of performing clustering in only the VNIR because, in the range of PRISMS (400-850nm), multiple different pigments can often be mixed in such a way that they may possess similar spectra which may have no, or very subtle, spectral differences which will likely not be distinguished during clustering. This subsection will briefly mention some of these incorrectly identified mixtures so that there is a greater understanding behind some of the pigment results displayed in the pigment mixture table within 3.5.

Yellow (Ochre, Orpiment, Organic, Gold) misidentification

Probably the most common form of misidentification is when clustering and grouping together thinly applied yellow-painted regions. While it was shown in chapter 2 to be expected for gold regions to become misclustered into certain clusters representing yellow pigments, the continued grouping highlighted that multiple different yellow pigments could be incorrectly clustered together due to there being only subtle differences in spectral shape and a quick transition to where the painting materials often became transparent (which meant that some absorption features would be much weaker toward the NIR). A good example of this can be seen in Fig 3.99, where a mixture already identified in the pigment ID section to be yellow ochre in Or 15932, could be incorrectly grouped into a cluster mainly representing a thin application of orpiment across the rest of the collection. As the ochre is not very thickly applied in this manuscript the absorption features are weak, resulting in spectra that can be clustered together within the colour difference and standard deviation range of spectra for light applications of orpiment, even though under closer inspection a subtle spectral difference can be identified.

In this scenario, the main cause of this misidentification is most likely due to the standard deviation range being larger due to an increased spectral contribution of the paper towards the NIR, which for many manuscripts tends to have an absorption feature at approximately 650nm, approximately the same location with which yellow ochre possesses its own features.

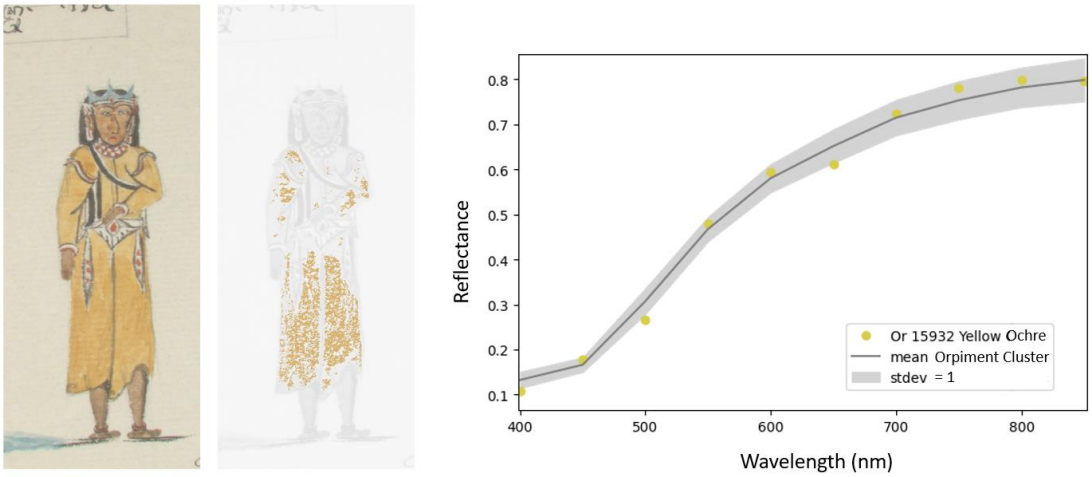

Fig.3.98 (left) RGB image and true colour cluster map of the yellow ochre which is clustered with the thinly applied orpiment. (Right) Graph showing how the yellow ochre mean cluster spectrum taken from Or 15932 fits within the mean cluster spectrum over the entire collection, due to the weak absorption features keeping all data points within 2 standard deviations of the mean.

This presents obvious problems, as by eye it is often quite clear from a spectrum of yellow ochre that its absorption features can be used to separate it from other yellow pigments. Such problems are quite typical of performing clustering in Euclidean distance, so some potential fixes could be to implement different distance metrics in the future, or perhaps to use alternative data, such as the derivative, to accentuate absorption features over intensity. However even when doing this one must be careful, as by changing one part of the data something else may likely be altered and begin to become misclustered instead. Using finer parameters in the algorithm may also work, but this is likely to result in creating many other clusters which may already be perfectly correct, thereby needlessly increasing the work required for grouping many other pigments and mixtures. A good alternative instead may therefore be to introduce brand new data, as implementing alternative spectral imaging techniques directly into the clustering may benefit the accuracy of the clusters in the same way that adding in colour could, where, for example, orpiment and yellow ochre for this cluster could be separated by their iron content in XRF as an addition to the VNIR. This idea is explored more in chapter 4, where a form of holistic multimodal clustering is investigated, however for the analysis of these results, the existence of point measurements was satisfactory for providing the information needed to distinguish the incorrectly grouped spectra in individual manuscripts and therefore properly identify different artistic materials collection-wide.

Green (Indigo, Yellow) Misidentification

Another common misclustering can be seen when comparing mixtures of indigo with different types of yellow. In a scenario where two different yellows are mixed with indigo, the resulting mixture can generate an identical VNIR spectrum which differs more with concentration differences than with the type of yellow used, as demonstrated for multiple KM mixtures in Fig 3.100).

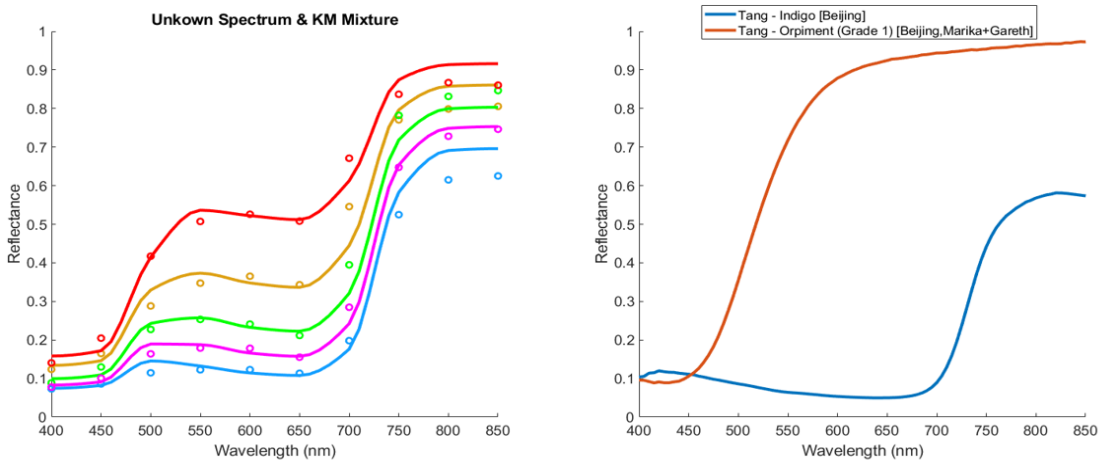

Fig.3.99 KM mixtures for different manuscripts which all fit well with a combination of paper and orpiment, regardless of the real yellow pigment used (organic or orpiment).

This means that some of the yellow pigments with more subtle spectral differences would minimally affect the green spectra, causing different mixtures to be clustered together based more on the concentration of any type of yellow rather than pinpointing exactly which yellow pigment is used. Again, as this is likely an intrinsic limitation of VNIR which could be solved by introducing new datasets, such as XRF or ER-FTIR mapping, directly into the clustering procedure, however, for this study, the numerous point measurements taken for different manuscripts were used to provide clarification.

Red (Vermillion, Vermillion + Red Lead, Vermillion) Misidentification

In a few of the Javanese manuscripts, it was noted that a red mixture likely consisting of vermillion and red lead (informed by XRF) was commonly used throughout many pages. For the VNIR spectra extracted from these red regions, many pixels with pure vermillion spectra can be clustered together with the potential vermillion/red lead mixture due to their similarity in spectrum and colour. This is another result which is to be expected considering that pure vermillion references mixed with paper in KM modelling can be fitted to the mixture just as well as the mixture of vermillion and red lead together (see Fig 3.101).

Therefore, during pigment mapping, many of the pixels representing these mixtures were consequently clustered and in turn grouped regardless of their differences in constituent materials. For these maps, however, XRF point measurements had been recorded, so segregation between manuscripts where additional red lead is used could still be performed for the pigment study.

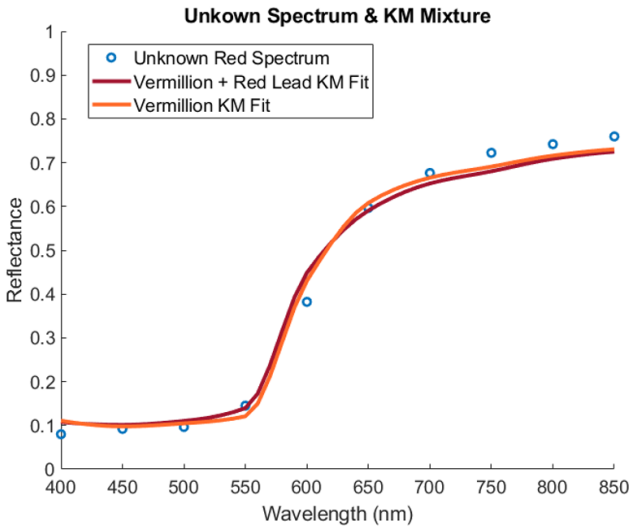

Fig.3.100 Graph showing how a mixture of vermillion and red lead fits just as well as pure vermillion for a red spectrum taken from MSS Jav 24

Other likely Misidentifications

With only these three examples it becomes clearer as to why some different mixtures may be misidentified during the clustering and grouping procedure, though with the point measurements, correct identification can still be performed. There are however other scenarios where the same problems can be encountered for other materials, such other examples include incorrectly clustering or grouping:

- Red dyes (with small absorption features) with vermillion.
- Different red dyes with one another
- Dark red with dark orange.
- Additional green mixtures (similar to indigo + orpiment problems)
- Different white colours and paper.

3.4 Collection-Wide Discussion & Conclusion

A convenient summary of the results found both through the point measurements and via clustering can be easily demonstrated in the form of a pigment mixture table, seen in Fig 3.102-3, where filled cells correspond to one of the pigments or pigment mixtures in each column appearing within the manuscript in each row. Cells are labelled with their potential origin location and those marked by asterisks are clustering results which cannot be confirmed with point measurements in any way.

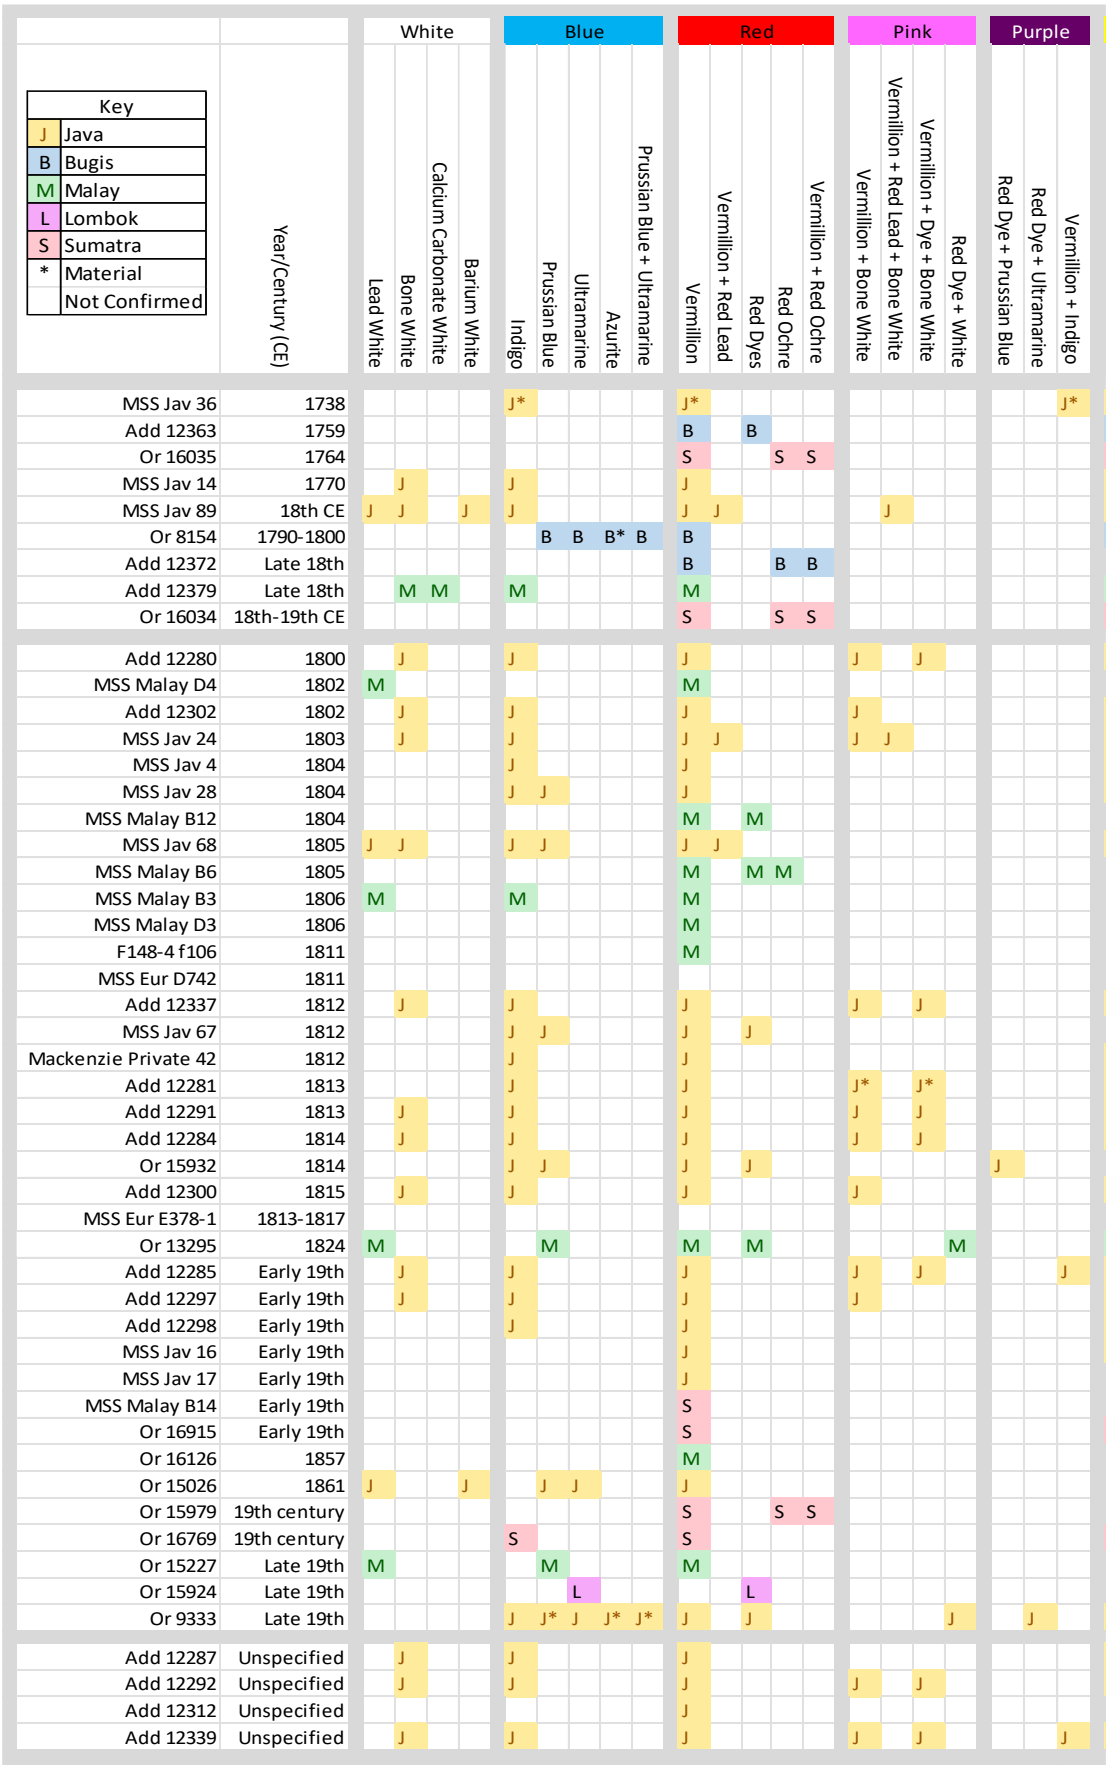

Key	
J	Java
B	Bugis
M	Malay
L	Lombok
S	Sumatra
*	Material
	Not Confirmed

	Year/Century (CE)	White				Blue					Red					Pink				Purple		
		Lead White	Bone White	Calcium Carbonate White	Barium White	Indigo	Prussian Blue	Ultramarine	Azurite	Prussian Blue + Ultramarine	Vermillion	Vermillion + Red Lead	Red Dyes	Red Ochre	Vermillion + Red Ochre	Vermillion + Bone White	Vermillion + Red Lead + Bone White	Vermillion + Dye + Bone White	Red Dye + White	Red Dye + Prussian Blue	Red Dye + Ultramarine	Vermillion + Indigo
MSS Jav 36	1738					J*					J*											J*
Add 12363	1759										B		B									
Or 16035	1764										S			S	S							
MSS Jav 14	1770		J			J					J											
MSS Jav 89	18th CE	J	J		J	J					J	J					J					
Or 8154	1790-1800						B	B	B*	B	B											
Add 12372	Late 18th										B			B	B							
Add 12379	Late 18th		M	M		M					M											
Or 16034	18th-19th CE										S			S	S							
Add 12280	1800		J			J					J					J		J				
MSS Malay D4	1802	M									M											
Add 12302	1802		J			J					J					J						
MSS Jav 24	1803		J			J					J	J				J	J					
MSS Jav 4	1804					J					J											
MSS Jav 28	1804					J	J				J											
MSS Malay B12	1804										M		M									
MSS Jav 68	1805	J	J			J	J				J	J										
MSS Malay B6	1805										M		M	M								
MSS Malay B3	1806	M				M					M											
MSS Malay D3	1806										M											
F148-4 f106	1811										M											
MSS Eur D742	1811																					
Add 12337	1812		J			J					J					J		J				
MSS Jav 67	1812					J	J				J		J									
Mackenzie Private 42	1812					J					J											
Add 12281	1813					J					J					J*		J*				
Add 12291	1813		J			J					J					J		J				
Add 12284	1814		J			J					J					J		J				
Or 15932	1814					J	J				J		J							J		
Add 12300	1815		J			J					J					J						
MSS Eur E378-1	1813-1817																					
Or 13295	1824	M					M				M		M						M			
Add 12285	Early 19th		J			J					J					J		J				J
Add 12297	Early 19th		J			J					J					J						
Add 12298	Early 19th					J					J											
MSS Jav 16	Early 19th										J											
MSS Jav 17	Early 19th										J											
MSS Malay B14	Early 19th										S											
Or 16915	Early 19th										S											
Or 16126	1857										M											
Or 15026	1861	J			J		J	J			J											
Or 15979	19th century										S			S	S							
Or 16769	19th century					S					S											
Or 15227	Late 19th	M					M				M											
Or 15924	Late 19th							L					L									
Or 9333	Late 19th					J	J*	J	J*	J*	J		J						J		J	
Add 12287	Unspecified		J			J					J											
Add 12292	Unspecified		J			J					J					J		J				
Add 12312	Unspecified										J											
Add 12339	Unspecified		J			J					J					J		J				J

Fig.3.101 LHS of the pigment table for the British Library's maritime Southeast Asian manuscript collection

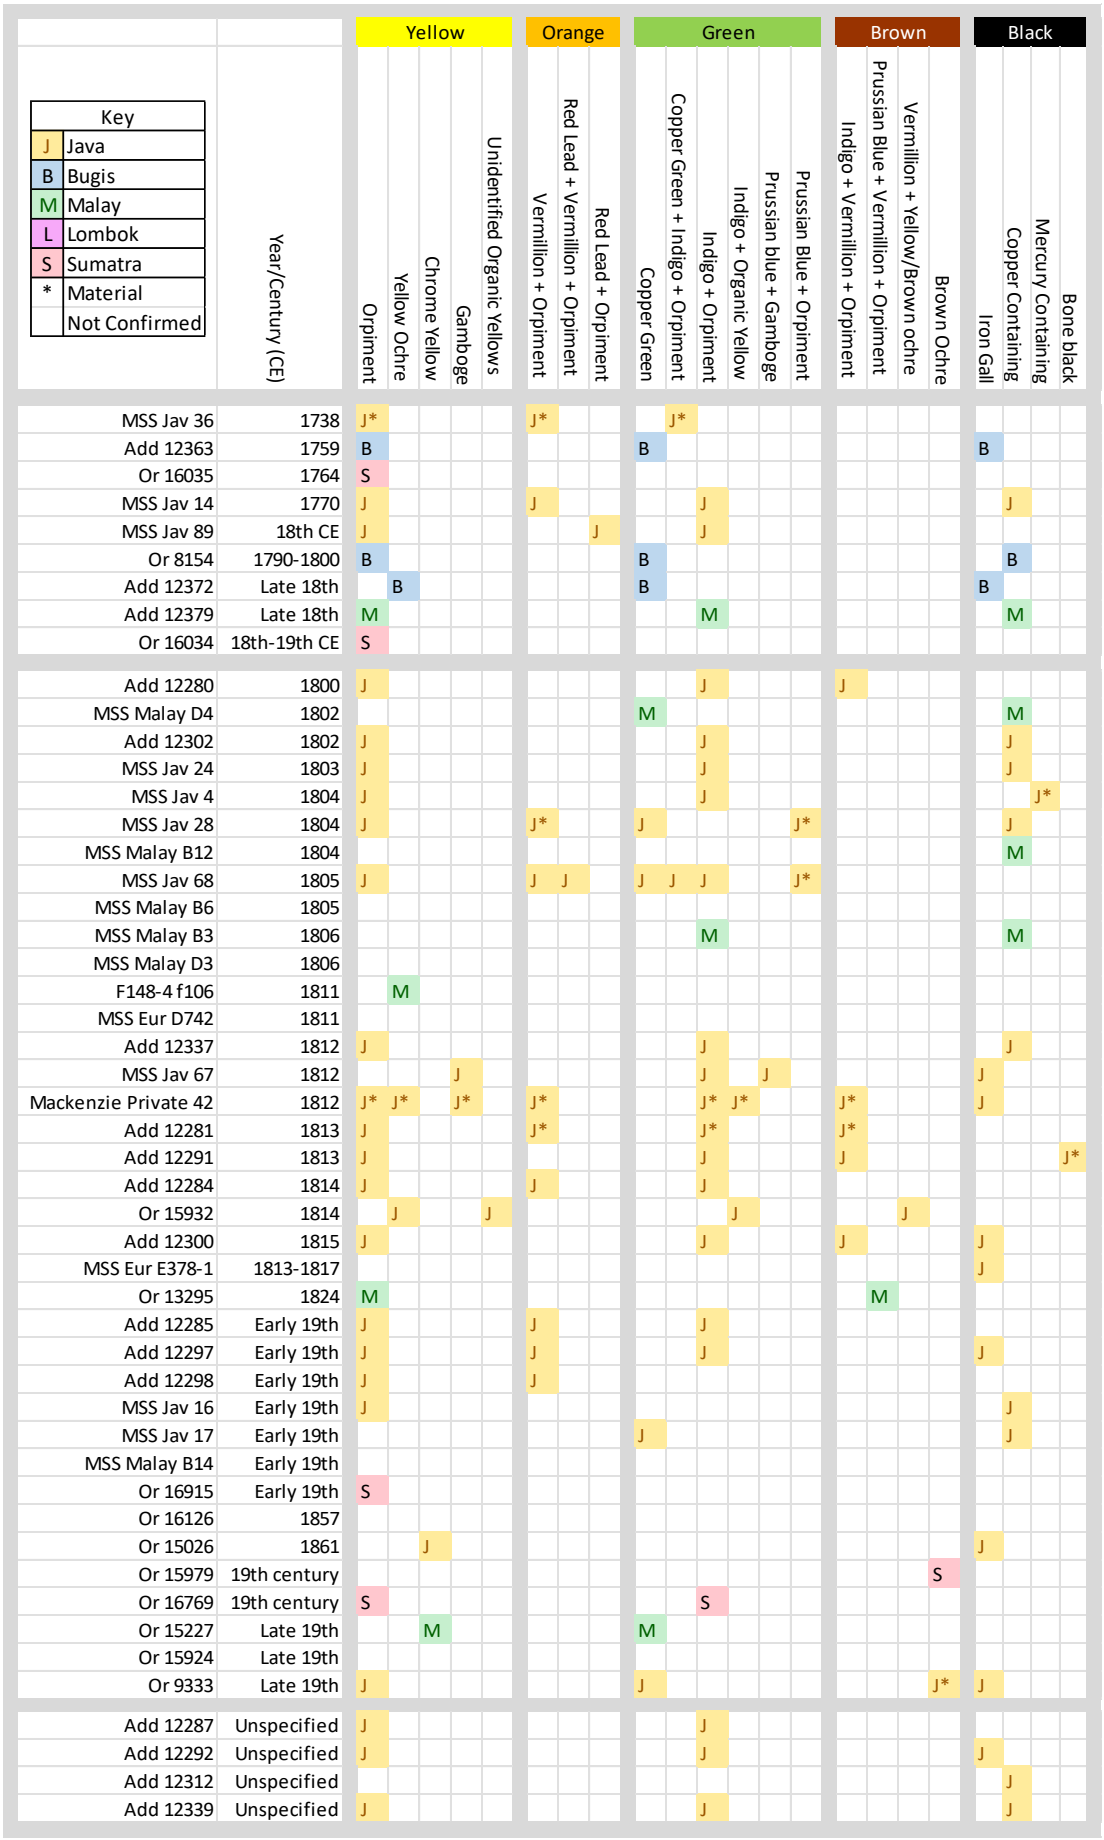

Key	
J	Java
B	Bugis
M	Malay
L	Lombok
S	Sumatra
*	Material
	Not Confirmed

		Yellow					Orange			Green					Brown				Black			
	Year/Century (CE)	Orpiment	Yellow Ochre	Chrome Yellow	Gamboge	Unidentified Organic Yellows	Vermillion + Orpiment	Red Lead + Vermillion + Orpiment	Red Lead + Orpiment	Copper Green	Copper Green + Indigo + Orpiment	Indigo + Orpiment	Indigo + Organic Yellow	Prussian blue + Gamboge	Prussian Blue + Orpiment	Indigo + Vermillion + Orpiment	Prussian Blue + Vermillion + Orpiment	Vermillion + Yellow/Brown ochre	Brown Ochre	Iron Gall	Copper Containing	Mercury Containing	Bone black
MSS Jav 36	1738	J*					J*				J*												
Add 12363	1759	B								B										B			
Or 16035	1764	S																					
MSS Jav 14	1770	J					J					J									J		
MSS Jav 89	18th CE	J							J			J											
Or 8154	1790-1800	B								B											B		
Add 12372	Late 18th		B							B										B			
Add 12379	Late 18th	M										M									M		
Or 16034	18th-19th CE	S																					
Add 12280	1800	J										J			J								
MSS Malay D4	1802									M											M		
Add 12302	1802	J										J									J		
MSS Jav 24	1803	J										J									J		
MSS Jav 4	1804	J										J										J*	
MSS Jav 28	1804	J					J*			J					J*						J		
MSS Malay B12	1804																				M		
MSS Jav 68	1805	J					J	J		J	J	J			J*								
MSS Malay B6	1805																						
MSS Malay B3	1806											M									M		
MSS Malay D3	1806																						
F148-4 f106	1811		M																				
MSS Eur D742	1811																						
Add 12337	1812	J										J									J		
MSS Jav 67	1812				J							J		J						J			
Mackenzie Private 42	1812	J*	J*		J*		J*					J*	J*			J*				J			
Add 12281	1813	J					J*					J*				J*							
Add 12291	1813	J										J				J							J*
Add 12284	1814	J					J					J											
Or 15932	1814		J			J							J					J					
Add 12300	1815	J										J				J				J			
MSS Eur E378-1	1813-1817																			J			
Or 13295	1824	M															M						
Add 12285	Early 19th	J					J					J											
Add 12297	Early 19th	J					J					J								J			
Add 12298	Early 19th	J					J																
MSS Jav 16	Early 19th	J																			J		
MSS Jav 17	Early 19th									J											J		
MSS Malay B14	Early 19th																						
Or 16915	Early 19th	S																					
Or 16126	1857																						
Or 15026	1861			J																J			
Or 15979	19th century																		S				
Or 16769	19th century	S										S											
Or 15227	Late 19th			M						M													
Or 15924	Late 19th																						
Or 9333	Late 19th	J								J									J*	J			
Add 12287	Unspecified	J										J											
Add 12292	Unspecified	J										J								J			
Add 12312	Unspecified																				J		
Add 12339	Unspecified	J										J									J		

Fig.3.102 RHS of the pigment table for the British Library's maritime Southeast Asian manuscript collection

With this table, the material identification results become much easier to interpret, and as can be seen many of the results found throughout the collection presented brand new findings as the study into the use of different pigments within maritime Southeast Asia had been uncommon prior to this research. Furthermore, as clustering was able to demonstrate the distribution of different pigment mixtures in large regions and over many items, the presence of different materials could be more confidently determined when paired with point measurement analysis such as XRF, Raman spectroscopy, and ER-FTIR. In fact, it was often shown that the pigment mixture maps produced after clustering and grouping VNIR spectra were only further supported by the extra XRF measurements regularly recorded for evaluating the clustering technique, where VNIR data was able to correctly identify material composition in most circumstances. It was clear therefore, that studying such a large collection in this way delivered great insight and efficiency that was invaluable in the study of artistic practises within maritime Southeast Asia, thereby providing the following opportunity to discuss the use of pigments over the entire collection more confidently.

With regards to the white pigments, it appears that most regions within maritime Southeast Asia still tend to use more traditional reserved white within their illumination, or just happen to have used no white for the items covered in this thesis. This could be seen in multiple manuscripts from Aceh, Lombok, Borneo, and Sulawesi, however for both the Javanese and Malay manuscripts, this does not appear to be the case. For most Javanese manuscripts, the most used white pigment appears to be one created from bone ash, and while its use seems logical due to Java having a history of using bone ash within the creation and painting of Wayang shadow puppets, it does come as a surprise as the most used white pigment generally found during the 18th-19th centuries would have been expected to be lead white. Another interesting discussion point here is that three of the four Javanese manuscripts (bottom section of the table) with unspecified dates possess bone white and other similar colours to different manuscripts, which may place their dates of creation at the early 19th century.

The Malay manuscripts on the other hand (along with a three Javanese examples) primarily used lead white in their illumination implying a greater willingness to use non-traditional white pigment than was seen in Java. Most of these manuscripts can be dated to the early 19th century, so it seems reasonable to expect they may use similar materials, however, it does appear that toward the latter half of the 19th century, barium white had started to be used within the archipelago, perhaps indicating increased connection with Europe, as indicated by its presence in Or 15026 (dated to 1861 though there is also evidence of barium white usage in MSS Jav 89, created before 1815).

Blue colourants could be found implemented in many manuscripts throughout the collection. Most Javanese manuscripts tend to use indigo, however, there are some which chronologically have the same provenance but use Prussian blue, likely suggesting greater European influence, especially if one considers that the manuscripts not using indigo also do not tend to use the more

traditional bone ash white. Malay manuscripts within the collection have a 50/50 split, with the earliest two dated items using indigo and the latest two Prussian blue, and they all (except for the Javanese influenced Add MS 12379) use lead white to lighten the blue colours. Ultramarine seems to be used primarily by only the latest manuscripts (after 1860), likely suggesting the use of artificial ultramarine and therefore, again, an increasing European influence towards the late 19th century. The only exception to this is Or 8154 (dated 1790-1800), which is unique in its use of both Prussian blue and ultramarine being mixed, which could suggest that ultramarine may have been applied later or that the natural variation is used. There is also some limited evidence of azurite being used within Or 9333 and Or 8154, however to be sure of this identification further analysis would be required.

For red pigments, the use of vermillion dominates throughout the collection, with it being present in all but one item where red is used, where in Or 15924 (a scroll from Lombok) madder lake was potentially likely to be used instead. The use of vermillion is no surprise, as during the 18th-19th centuries it was the most used red pigment in Asia, however, it is interesting to note that in many scenarios it is not the only red pigment used. In multiple Javanese manuscripts, the mixture of red lead and vermillion was potentially likely, and in multiple other manuscripts, red dyes are often implemented used as standalone colourants or are mixed with white, however there was also limited evidence suggesting a potential mixture with vermillion. Red dyes could be mixed with white in both Malay and a late Javanese manuscript, Or 13295 and Or 9333 respectively. They could also be found applied straight onto paper in numerous Malay manuscripts, a single Bugis manuscript (Add MS 12363), and a few Javanese manuscripts, though for these manuscripts from Java the red dye only occurred for the more westernised third artistic style mentioned in section 1. For the most part, red insect-based dyes were normally detected through the comparison of fine absorption features, however there could also be plant-based red dyes in areas where PRISMS was the only VNIR data collected. It is also important to note the additional use of red ochre in some manuscripts, with it appearing most commonly in Sumatra, in addition to one Malay (MSS Malay B6) and one Bugis manuscript (Add MS 12372).

When considering the use of yellow throughout the collection, the most abundantly used pigment is orpiment, in which evidence could be found for it to exist in almost every manuscript regardless of provenance. It was commonly used in Acehnese, Javanese, Bugis and a few Malay manuscripts, with other yellows seemingly used only seven times throughout the rest of the collection. The other yellows found included yellow ochre, found in a select few manuscripts, chrome yellow (found in two later manuscripts from 1861-late 19th century), and organic yellows, with one confirmed to be gamboge in MSS Jav 67 (dated to 1812). The use of ochre is no surprise as it has been used in general for thousands of years, however, the presence of chrome yellow likely suggests further European influence towards the end of the 19th century, and gamboge

perhaps an influence from mainland southeast Asia or China, where it was also used in the 19th century (Eremin et al., 2008; French & Monaghan, 2018)

Throughout the collection green is also used extensively, with the most abundant being a result of the mixture of indigo and orpiment, commonly seen within the Javanese manuscripts, however the same mixture can also be found in two Malay manuscripts and a single Sumatran manuscript. In addition to this mixture however can be found numerous others, such as a mixture of Prussian blue with gamboge (Hooker's green), indigo with organic yellow, and potentially Prussian blue with orpiment. Furthermore, there are also use of other copper-based greens in numerous Bugis, Malay, and Javanese manuscripts. The use of indigo and orpiment is typical of many cultures and had already been confirmed in a previous study for Javanese manuscripts by Burgio et al., (1999), however its use at the scale demonstrated by this study clearly indicates its importance in maritime Southeast Asia. The use of Prussian blue within green mixtures is likely due to European influence during the 19th century, as the manuscripts which use Prussian blue all tend to have similar dates. However, the use of copper-based green has an ambiguous source depending on different regions, as the Bugis manuscripts used copper-based greens very early (mid-late 18th century) but others, such as the Malay manuscript Or 15227, were produced towards the late 19th century.

For most of the other colours, it seems that the common practice was to achieve different shades and hues by mixing primary colourants seen elsewhere in a manuscript. This can include mixing white and red or blue to make pinks or different shades of blue, mixing reds and blue to create purples, red and yellow for orange, or many different colours to form a brown or others. This is an interesting result as it implies that many of the artists may have known how to properly mix different artistic materials to gain different shades and colours, a practise not necessarily guaranteed at the time. The only exception this for brown colours was where brown ochre could be likely be found on the Sumatran manuscript Or 15979 (dated to 19th century), though it is unknown if this was intentional or exists now as a degradation product of the red ochre originally applied. Finally, this holistic study was also able to provide insight into the inks used throughout numerous manuscripts in various parts of maritime Southeast Asia, where 24 manuscripts (around half the collection) were confirmed to use either iron gall ink or a copper-containing ink (which could correspond to numerous inks such as black copper (II) oxide or carbon black, though to perform identification of the inks at a grand scale it is likely further Raman spectroscopy would be required. There was suggestive XRF evidence in favour of bone black in Add MS 12291, though XRF alone is not enough to confirm this, and seemingly a mixture of vermillion and black in MSS Jav 4.

Therefore, the study overall appeared to show an adoption of new pigments towards the mid-late 19th century, most likely due to European influence where materials such as barium white and chrome yellow began to appear. But with the collection skewed towards the early-19th century, it is difficult to find any consistent trends with how pigment usage may have changed with time.

One factor that can be determined however is that within the early 19th century, most manuscript illumination involved same materials, as can be seen by the abundance of the same mixtures appearing for different blues, yellows, reds, and greens, especially in Java, where most of the collection is from.

Chapter 4

Holistic Multimodal Clustering of Export Paintings

4.1 Introduction

Spectral Imaging techniques of different varieties are commonly used in cultural heritage studies and can be useful for researchers and curators as they can allow for the fast non-invasive collection of large amounts of spectral data. When analysing paintings, a broad range of spectral features can be encountered for different techniques due to the variation of artistic materials and styles used. In these studies, some commonly implemented spectral imaging techniques can include VNIR spectral imaging (Cucci et al., 2016; Liang, 2012; Picollo et al., 2020), XRF mapping (Alfeld et al., 2011; Caliri et al., 2021; Mazzinghi et al., 2021; Saverwyns et al., 2018), and ER-FTIR mapping (Legrand et al., 2014). Each technique used can produce complex results which require considerable effort to analyse individually, and when considering that many painting studies require a holistic approach to fully analyse materials, every technique added into the analysis can exponentially increase the complexity of the data. Consequently, material analysis for paintings can be time consuming and, in some scenarios, certain features may be missed if a researcher interprets where to extract spectra through observations made only by eye. In the previous chapters clustering was shown to provide a convenient solution to tackling complex spectral imaging data, though this was only performed for the VNIR modality of a holistic study. In the event where multiple spectral imaging techniques are required to be performed on a single painting, it is desirable to have a more automated time-efficient approach for each method. In these scenarios, the clustering technique outlined in chapter 2 could be used to reduce the total number of spectra and create cluster maps for each spectral imaging modality, thereby making the analytical process less intensive while also reducing the chances of missing important information. Furthermore, while clustering the spectral imaging data produced from each technique would provide advantages over analysing multimodal data manually, clustering approaches do not have to be limited to one spectral domain at a time. Instead, with the correct pre-processing and data treatment, clustering can be performed across all modalities simultaneously, allowing for complementary information to be used to automatically produce cluster maps which have features created from all techniques, which in turn may reveal information about a painting which would otherwise be difficult to identify. While many approaches have been suggested and developed for this type of data fusion in many applications (Lahat et al., 2015) including remote sensing (Chang & Bai, 2018; Ghamisi et al., 2019), its use in cultural heritage applications has not been thoroughly explored, and in most scenarios where it has been implemented, priori knowledge and direct interaction with the data in feature space can be required such as in Alfeld et al., (2018). This implies that in depth manual research is still

required to some degree before implementing any machine learning approaches, however in practical terms the multimodal study of paintings would be best executed when using unsupervised clustering at the beginning of the study. This chapter investigates a novel approach for performing material identification by examining how automated clustering can be implemented at the beginning of an analysis workflow for multiple complementary spectral imaging techniques. Exploration into how the results from each modality can be unified into a single cluster map is also performed to allow every cluster a unique spectrum in all spectral domains. Human interaction with the data was not required until the final stages, where a reduced number of unique clusters were assigned into groups for performing pigment identification. Two different studies were performed on export paintings from Peru and China, where different complementary information was exposed to a multimodal clustering algorithm which can allow for more accurate and efficient confirmation of pigments and pigment mixtures.

4.2 Holistic Multimodal Clustering of Peruvian Export Paintings

To perform more automated holistic studies of different paper-based paintings, clustering using techniques of different modalities was required where each would provide its own portion of complementary data required to complete a full pigment study. The first dataset consisted of a Peruvian export watercolour painting with a relatively simple palette which included various red, yellow, brown, blue, green, and black pigments, with a reserved white. In the analysis performed for the British Library's maritime southeast Asian collection in chapter 3, it was shown that VNIR spectral imaging could be used alongside XRF to aid in the identification of artistic materials. The complementarity of the two techniques was useful when attempting to identify pigments with similar VNIR spectra but different elemental composition, such as in distinguishing different yellow paints. Furthermore, XRF was useful for guiding the interpretation and analysis of results by directing many of the mixtures used to identify pigments using the Kubelka-Munk model. VNIR and XRF spectral imaging data were recorded for the Peruvian export painting, P80001-8, with the intention of performing a manual holistic study of the pigments used in the painting. However, after the success seen in combining colour and reflectance information together in previous chapters, the procedure of automated clustering and grouping was seen as a good candidate for performing multimodal clustering across completely different techniques, so that the cross comparison of artistic materials and their application could be performed more efficiently.

4.2.1 VNIR Spectral Imaging

VNIR and colour data was recorded using the PRISMS (Liang et al., 2014) system previously mentioned in chapters 2 and 3, however each pixel possessed a VNIR spectrum ranging from 400-880nm for this study, separated into nine 40nm bandwidth filters at 50nm intervals from 400-800nm and a final filter at 880nm with a bandwidth of 70nm.

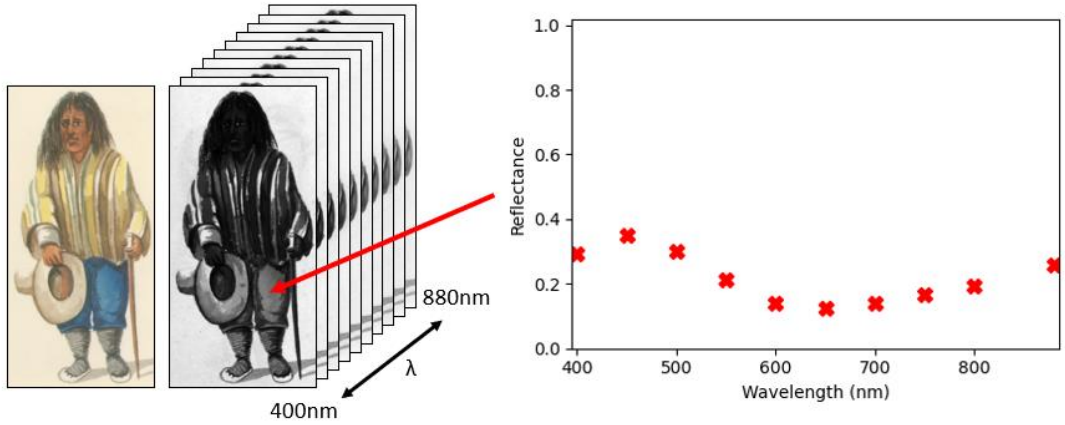

Fig.4.1 Left: Representation of the 10 PRISMS spectral imaging channels placed into a 3D spectral imaging data cube. Right: An example reflectance spectrum representing a blue area on the painting.

4.2.2 XRF Spectral Imaging & Pre-Processing

Ma-XRF data was recorded for using a Bruker M6 "Jetstream" XRF scanner (Bruker, 2023). The M6 used a spot size of 230µm for data collection, with a sampling step of 200µm and a dwell time of 25ms. Voltage and current were set to 50kV and 600mA, respectively. MA-XRF data cubes consisted of a 2D image where each pixel possessed a 4096-channel ranging 0.35-40keV. Therefore, the XRF spectral data cube had a total size of 276x570x4096 pixels, a total of 157320 spectra with 4096 channels each (Fig. 4.2).

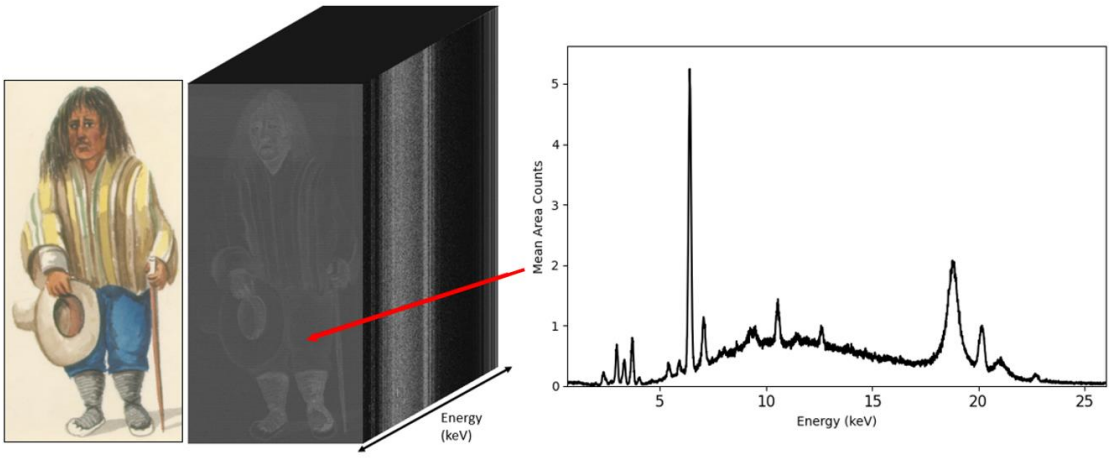

Fig.4.2 A representation of the XRF spectral imaging data cube collected by the Bruker M6, with an example spectrum extracted from the blue area of the painting.

Pre-Processing

The clustering of MA-XRF Spectral Imaging data using SOMs had already been demonstrated by the ISAAC research group for a painting in this Peruvian collection before in (Kogou et al., 2021). The process in that investigation involved a prior pre-treatment step where a sum image was recorded for each detected XRF spectral line by adding together the 5 most intense channels around the top of each peak after applying a spatial median filter. This was useful in this investigation as it allowed for dimensionality reduction to be performed on the XRF data while maintaining the core information useful for the clustering of separate elemental mixtures, it also in turn would have increased the accuracy of SOM, due to better SNR and because in general for any clustering using Euclidean distance, fewer dimensions is preferable. There were however some concerns which prompted further research into how clustering may react with different pre-processing approaches, especially those which can be automated by a computer to allow for the detection of certain elements which a human may miss, such as in the detection of trace elements. In addition, the use of Euclidean distance as a standard metric for SOM based clustering meant that the use of sum images may result in unfairly weighted clustering results, where clusters prioritise small relative differences in strong XRF peaks over large relative differences in weak peaks. To counteract such issues, a more automated XRF pre-processing algorithm was designed to transform the full XRF spectral image into a smaller set of elemental spectral bands which could be used for clustering. The pre-processing of XRF data involved:

1. Calculating and subtracting the XRF background continuum from the spectral cube mean spectrum.
2. Finding all XRF peaks in the continuum subtracted spectrum.
3. Creating elemental spectral bands by merging the energy channels within the FWHM of each peak into mean images before applying a 3x3 spatial median filter.

Continuum Subtraction

The removal of Bremsstrahlung background radiation within XRF data is an important part of processing as it can make the procedure of identifying elements, and therefore image channels, much easier if there is no underlying background feature to skew the heights of different energy peaks. Therefore, a reliable and accurate method of performing continuum subtraction on XRF mapping data was applied to be able to separate the background from the channels required for clustering. The technique decided upon to perform the subtraction was peak stripping, a commonly used method of estimating the background continuum which compares the counts in an XRF channel against the mean of its neighbours to remove structures, such as peaks or noise, in a spectrum. The process is covered in detail in (Van Grieken & Eds, 2002) the discussion continues further to cover how the application of smoothing algorithms and large numbers of

iterations can change the efficiency of the algorithm, where some options may provide more accurate results. As it is unsure of how well these different peak stripping methods may be applied onto data recorded by the Bruker M6, a short investigation was carried out to investigate three different approaches: Peak Stripping, Square Root Peak Stripping, and Log peak stripping.

This investigation for P80001-8 was performed by comparing an ideal 'reference continuum' calculated by a human researcher against the automated peak stripping results. A value for the standard deviation, σ, between the reference continuum and the peak stripping continuum was calculated over increasing numbers of iterations. Standard deviation was derived from the difference between each respective peak stripping approach and the manually selected continuum for multiple iterations, of which the calculations are as follows:

$$f(x)_i = m(x) - p(x)_i \quad , \quad \sigma = \sqrt{\frac{\sum_{i=1}^{N} \left(f(x) - \overline{f(x)} \right)^2}{N}}$$

Where $m(x)_i$ is the manual continuum subtracted spectrum, $p(x)_i$ is the peak-stripping continuum subtracted spectrum for iteration i and N is the number of XRF channels used. The changes in standard deviation were recorded for iterations between 1-10,000, with the standard deviation decreasing rapidly before levelling off and reaching a point of convergence, as can be seen in Fig. 4.3. The results from this test clearly show that the square root and log peak stripping approaches both offer an improvement over the ordinary peak stripping method. As it is difficult to distinguish which of the two improvements are best, a mean of the standard deviation was calculated for all the datapoints, allowing for a quantitative comparison which finally showed log peak stripping to be the best technique, though only by a small amount, as seen in the table in Fig. 4.3.

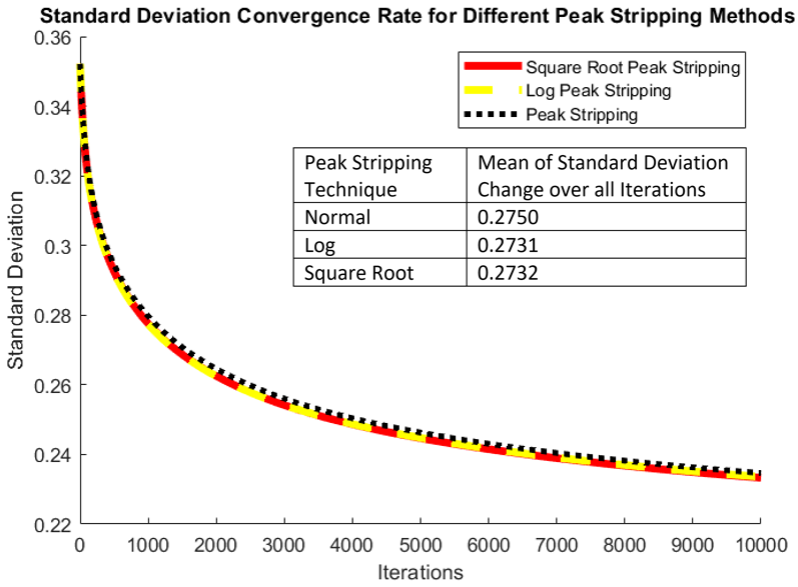

Fig.4.3 Graph showing the standard deviation convergence of three peak stripping techniques used for performing continuum subtraction. A table is included within the graph which shows the Log peak stripping to have the smallest deviation.

Regardless, for all techniques the curves clearly show that the standard deviation always decreases as iterations increase before reaching a position of convergence, suggesting that the peak stripping method will always eventually reach a point where no significantly better fitting is made to a continuum that is close to a human fitted reference continuum. Furthermore, as the process of performing log peak stripping is not too computationally expensive for a single mean XRF spectrum, it is never a problem for high numbers of iterations to be used. After the testing of the techniques, the process was applied onto the mean XRF spectrum to create a new continuum subtracted spectrum (see Fig. 4.4) so that peak fitting could be performed more easily.

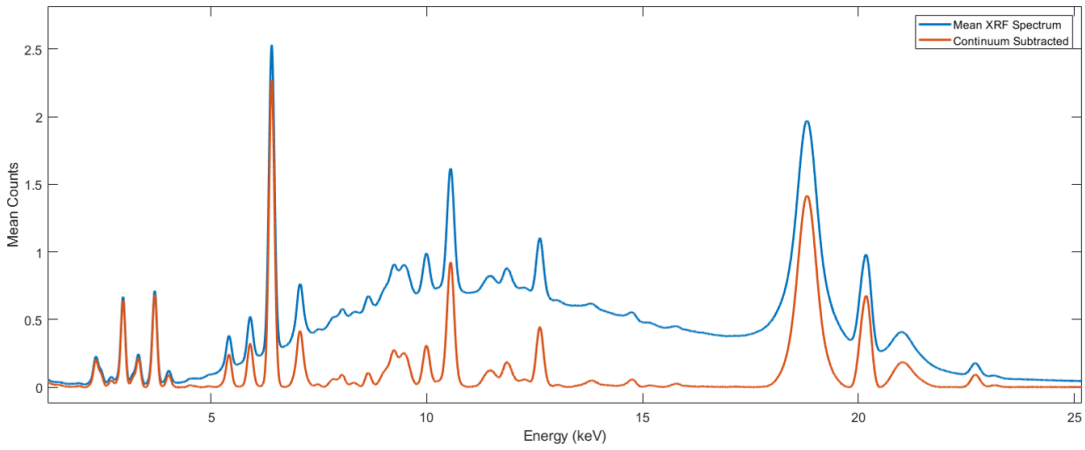

Fig.4.4 Graph showing the original mean spectral imaging XRF spectrum (in blue) compared to the continuum subtracted XRF spectrum (in orange), after log peak stripping.

Peak Fitting & Spectral Energy Band Creation

After the removal of the background continuum, different peaks for the entire XRF cube are found by detecting turning points and troughs in the first and second derivatives of the spectrum respectively. The process finds potential peak positions, checks the surrounding data points if a negative polynomial can be fitted (acting as the top of the peak), and stores peak positions for later processing, allowing for the detection of all energy peaks including trace elements, which may otherwise be missed if using XRF data with low signal to noise spectra or the more commonly used single element maps. After the peaks are found, gaussian curves are fitted to the surrounding data about each peak position and a FWHM is calculated so that the number of channels used to create new XRF energy bands scales with the width of the of the peak. The channels encompassed by the FWHM for each peak are subsequently combined into a single mean XRF energy channel, which can then be used in clustering in the same way a single PRISMS channel would be in the VNIR. In total for P80001-8, 29 individual XRF spectral channels were created from regions which possessed clear peaks (as seen in Fig. 4.5), covering many different chemical elements. For one of the channels at ≈7.82keV, the element could not be defined as it may instead correspond to an instrumental sum or escape peak as mentioned in chapter 1, though when attempting to calculate which, no strong signals corresponded well. Following this, the 29 channels were combined into a single XRF spectral imaging data cube, which was then applied to the automated clustering code covered in chapter 2, the results of which can be seen in section 4.2.5 of this chapter.

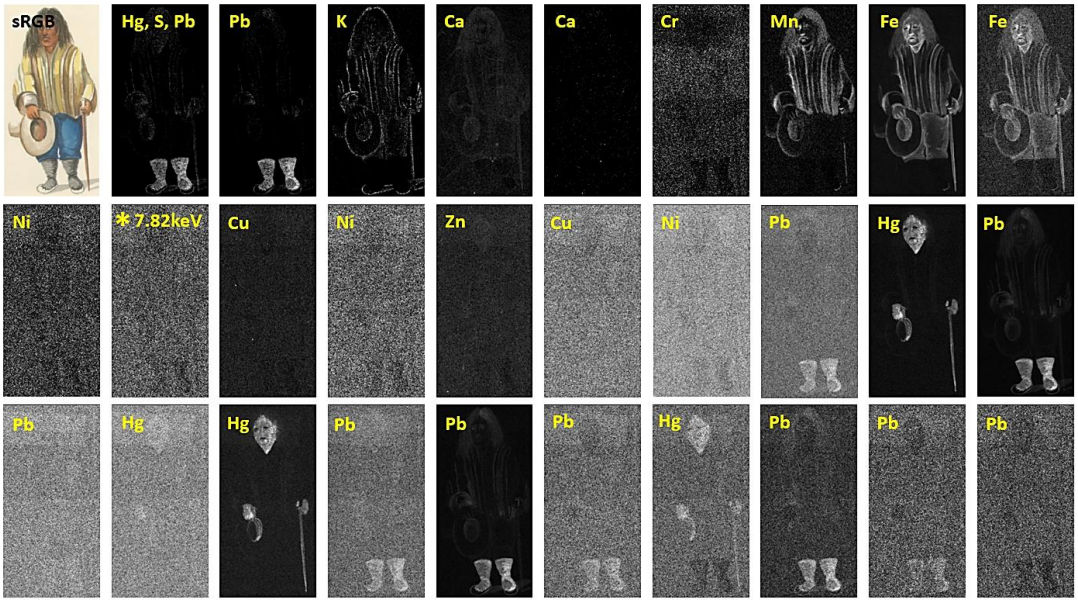

Fig.4.5 Collection of images representing the 29 different energy channels created after performing XRF pre-processing. Each channel is associated with an elemental peak except for E=7.82keV, where a peak was found but could not be associated with any known element, suggesting it may be instrumental.

4.2.3 Image Registration

Prior to performing multimodal clustering, VNIR, Colour, and ma-XRF spectral imaging data first needs to be registered so that the pixel size and resolution is comparable between the techniques. As the resolution for the ma-XRF was lower than other techniques, all VNIR images were scaled down to the same resolution as the ma-XRF maps. This was done by manually selecting common feature points between the sRGB image and mean ma-XRF image and using an affine transformation to perform a geometric transformation to the data that was also applied to the PRISMS spectral reflectance data and CIELAB data so that clustering could be performed. This registration was performed in MATLAB using the 'fitgeotrans' function, which is closely related to two papers, notably Goshtasby, (1986, 1988). A brief diagram illustrating the results of the registration along with the anchor points chosen can be seen in Fig. 4.6.

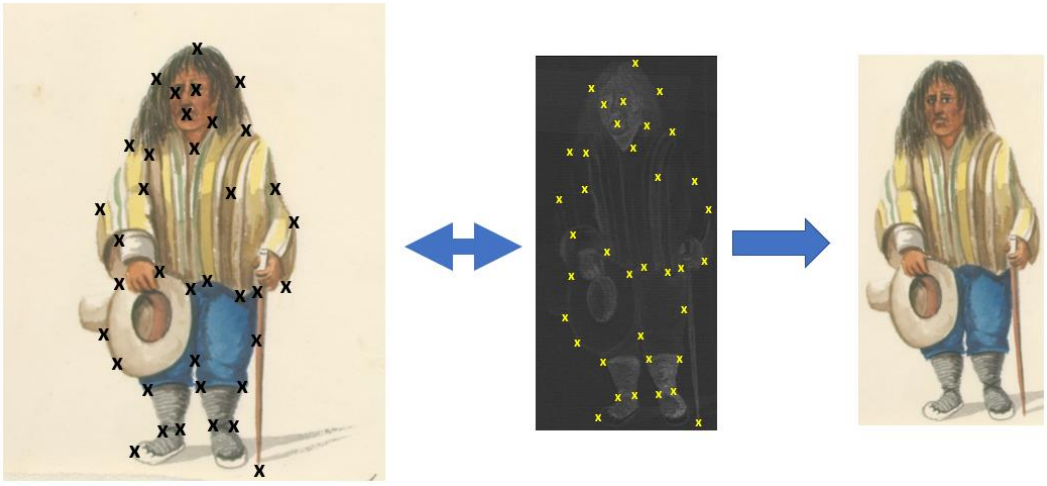

Fig.4.6 Representation of the image registration process performed to transform VNIR spectral and colour data (Left) down to the same resolution and pixel size as XRF (Middle). Crosses mark the anchor points used for the geometrical transformation of data, with a final RGB image showing the transformed data (Right).

4.2.4 VNIR Clustering Results

Following the registration of the VNIR spectral imaging data to the resolution and pixel size of the XRF data, the method outlined in chapter 2 was used to perform automated clustering for the P80001-8. For the clustering, parameters of colour difference, ΔE^*_{ab}=5, and standard deviation range, σ_{VNIR} = 2.5, were used as they were deemed adequate to ensure that the VNIR clustering was performed without any misclustering of colour or spectra. After undergoing filtering to remove redundant clusters, a total of 42 unique clusters were found over the painting, accounting for differences in pigment concentration and layering etc. A false colour cluster map generated prior to filtering can be seen next to the sRGB of the painting in Fig. 4.7, where, as expected, many redundant clusters appeared to exist in large numbers due to edge effects caused by artefacts resulting from the change in resolution during image registration. After performing clustering on

the VNIR data, the cluster maps produced can be used as inputs into the multimodal clustering technique, which is covered in detail in section 4.2.7.

Fig.4.7 Images showing the RGB image (left) of P80001-8 next to the false colour cluster map of the VNIR spectral data (Right) created after performing automated clustering.

Fig.4.8 Images showing the RGB image (left) of P80001-8 next to the false colour cluster map of the XRF spectral data (Right) created after performing automated clustering. Note: That the paper and yellow/green regions are clustered together due to their spectral similarity in the XRF.

4.2.5 XRF Clustering Results

Clustering using the 29-band pre-processed XRF spectral imaging data cube produced a total of 72 unique clusters, possessing many variations in intensities for XRF peaks in different combinations. For this clustering, a standard deviation for XRF, σ_{XRF}, of 2.5 was used, as this allowed for a set of clusters to be made where any smaller numbers of standard deviation did not significantly change any clustering results. Again, a false colour cluster map showing the clustering results prior to filtering can be seen next to the sRGB image for P80001-8 (Fig. 4.8). By observing this cluster map next to the sRGB image, the limitations of relying on only XRF alone to perform clustering can be seen clearly. Many pigments used in different paintings or other artistic works use only organic materials, which XRF would be unable to distinguish between as the detection limit for chemical elements typically only reaches as low as Silicon (Si) in air and does not cover the common components for organic materials such as oxygen, carbon or hydrogen. This is clearly demonstrated in the Peruvian painting where the background paper has been clustered together with the yellow and green parts of the man's shirt, which when observing the XRF spectra (as shown in Fig. 4.9) is clearly due to the two different coloured regions having the same elemental composition, even though a clear difference can be seen in the colour images.

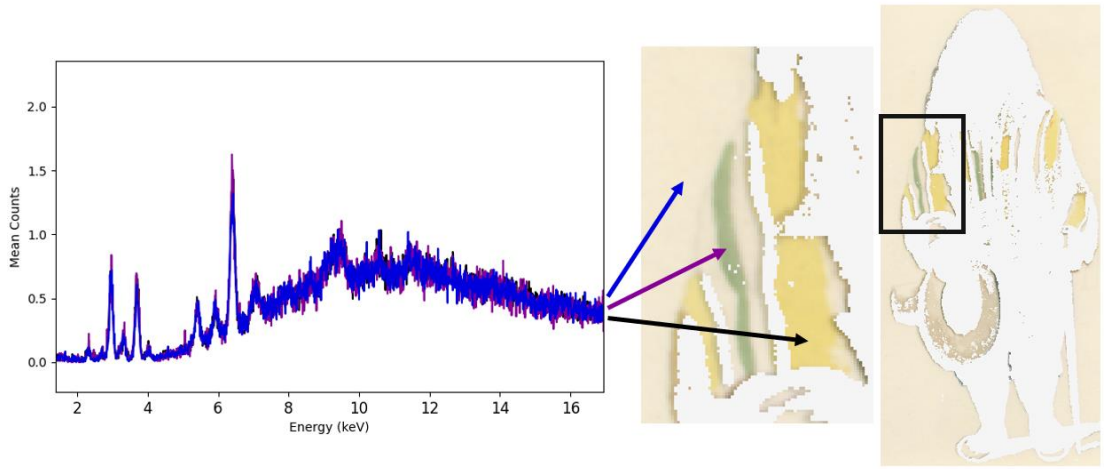

Fig.4.9 True colour cluster map of the yellow, green, and paper regions (right) which have been clustered together due to possessing the same XRF spectrum. Spectra for all three regions can be seen in the graph on the left-hand side of the figure.

Like with the VNIR clustering results outlined during chapter 3, using spectral imaging in only a single domain limits the ability of performing clustering for the purposes of pigment identification. This further justifies our need for multimodal clustering, especially when we consider that the yellow, green, and paper regions can all be clearly separated in the VNIR, as was demonstrated through the false colour map seen in Fig. 4.7.

4.2.6 Multimodal Clustering Procedure

To perform multimodal clustering, each individual technique first must be clustered individually, to produce output cluster maps which can subsequently be used as a new input. The basic idea for the process is to perform repeated clustering using the overlaying method, as mentioned in chapter 2, but using all modalities of data instead of just VNIR reflectance and CIELAB. As the different spectral imaging data cubes produced from each modality were registered to each other before being clustered individually, the different cluster maps can be overlaid easily and merged to produce a unique set of clusters ranging from 1 to N where every cluster should possess its own unique CIELAB L*a*b* vector, VNIR reflectance spectrum, and XRF energy spectrum, as demonstrated by the diagram in Fig. 4.10.

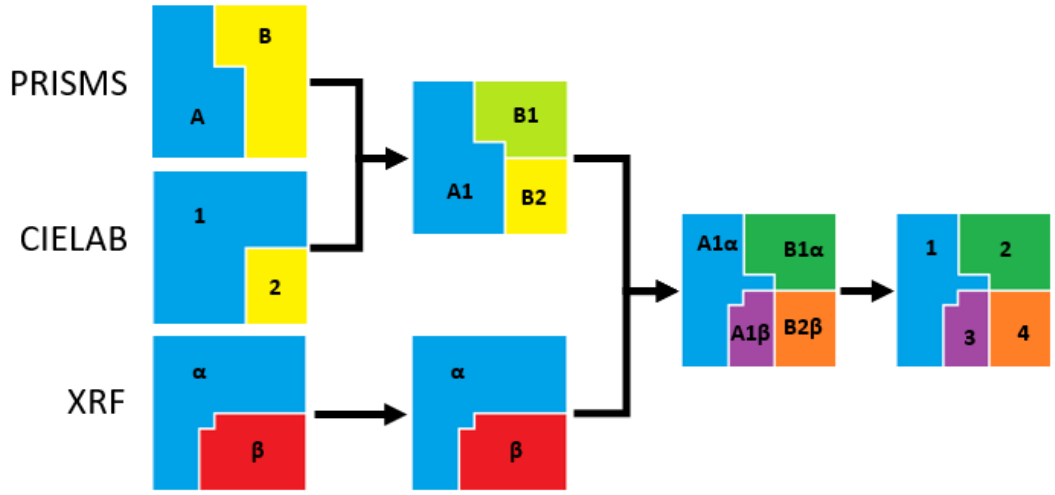

Fig.4.10 Basic graphical representation of the map overlay clustering method originally introduced in chapter 2, however this time implemented for automated clustering in three different techniques: PRISMS VNIR, CIELAB, and XRF.

After map overlay repeated clustering has been performed covering all three modalities, binary median filtering and hierarchical merging are performed once again to ensure that redundant clusters with low populations or sparse distributions are removed, and that everything is properly merged within the user defined parameters of σ_{XRF}, ΔE^*_{ab}, and σ_{VNIR}. For performing multimodal merging and removal, it is generally wise for the same user parameters to be used as were for each individual technique, however it has been seen that with the complementarity of each technique creating more refined clusters, larger parameters can sometimes be used without sacrificing accuracy. Finally, after hierarchical merging has been completed, the output data is produced, which again includes cluster maps, false colour cluster maps, and the statistical data for all techniques which is required for analysis and further grouping.

This full process of multimodal clustering is in theory able to accept as many different modalities as possible, and only has requirements of accurate image registration across techniques

beforehand. Therefore, in a scenario where many more techniques may be required for a full holistic analysis of a painting, such as in section 3, this methodology should continue to work, so long as each individual modality has been clustered correctly in the prior stages.

4.2.7 Multimodal Clustering Results

The produced XRF and Colour/VNIR cluster maps underwent the multimodal clustering procedure using the same parameters as were used for each individual modality: where σ_{XRF}=2.5, σ_{VNIR}=2.5, and ΔE^*_{ab}=5. After the process was completed the total number of clusters found within the multimodal cluster map was 114. A false colour representation of the 114 clusters can be seen in Fig. 4.11, with the removed redundant clusters coloured in black.

Fig.4.11 Images showing the RGB (left) of P80001-8 along with the resulting false colour cluster map produced after performing multimodal clustering, with redundant clusters removed and shown in black (right).

As expected, the final multimodal cluster map has contributions from both the XRF and VNIR, with each complementing the other in different ways. One such example showing how the complementarity of the techniques provides benefits over XRF alone is by monitoring the cluster which previously had joined the background paper with the yellow and green spectra on the shirt. Now, with the application of multimodal clustering, the separation of the clusters is performed due to adding the reflectance and colour information, as clearly shown in the false colour image seen in Fig. 4.10. Under close examination there are also other regions in which the same benefits to XRF can be found, such as for the blue trousers or the shaded parts of the hat, where more

correct clusters are formed due to the contribution of the VNIR spectral imaging. These improvements are almost expected, as differences in colour can easily be identified by eye, however it does provide a greater level of convenience and confidence in results to have all available information confirming the same distribution of materials represented in a single image.

Multimodal Clustering of Paint Layers

The full advantages of multimodal clustering however are better demonstrated when analysing clusters in the other direction; observing what additions XRF can provide to VNIR reflectance. One such example is the separation of cluster regions on the man's brown hair in the painting, where even though the spectral reflectance and colour are both identical, a separation of clusters exists due to the changes in the layering of the paints. These changes in paint layering can be identified in the XRF spectral data cube for the separately clustered brown region due to the presence of Mercury (Hg), as shown in Fig. 4.12, suggesting that a reddish pigment used for the skin, most likely vermillion, exists under only a small part of the brown hair. This presence of vermillion cannot be detected in the VNIR alone as demonstrated by the almost identical VNIR spectra, showing that such a result is only possible through the multimodal clustering of the two techniques.

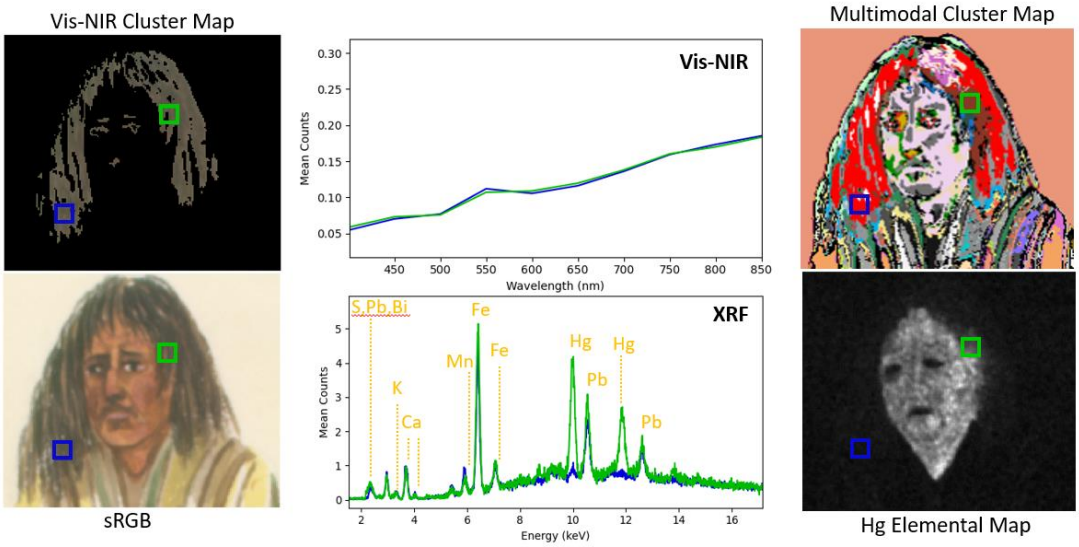

Fig.4.12 Representation of multimodal clustering identifying paint layers by incorporating XRF information into the PRISMS VNIR results.

Signal-Noise Ratio Improvements

In general, automated clustering provides mean spectra for different regions which will be of a much higher signal to noise ratio (SNR) than either a single pixel or even a small group of pixels which could be manually extracted from a spectral imaging data cube for analysis. This is obviously an improvement in general for performing any form of analysis, however, an additional

advantage of this for performing pigment identification is that small, subtle spectral features in low SNR regions can be revealed when at both the pixel level and in the individual elemental maps they would not be discernible. The multimodal clustering results illustrate this happening in multiple areas over the painting where the original pixel level XRF signal is too weak to be able to clearly separate small changes in elemental composition.

One good example is for green parts of the shirt which in XRF clustering alone were automatically grouped together with the background paper and organic yellow. For this region, manually extracting groups of pixels clustered in only the XRF showed that the spectra were identical, as already seen in Fig. 4.8, and it was therefore concluded that the clustering had performed well in joining these regions together. However, after performing multimodal clustering, the resulting higher SNR spectrum for the green parts of the shirt, which is only achievable through the contributions of the VNIR data, revealed a small difference in elemental composition between the green and the paper which was not discernible previously. This can be demonstrated clearly in Fig. 4.13, which revealed a more intense peak for iron (Fe) which, prior to multimodal clustering, could not be detected.

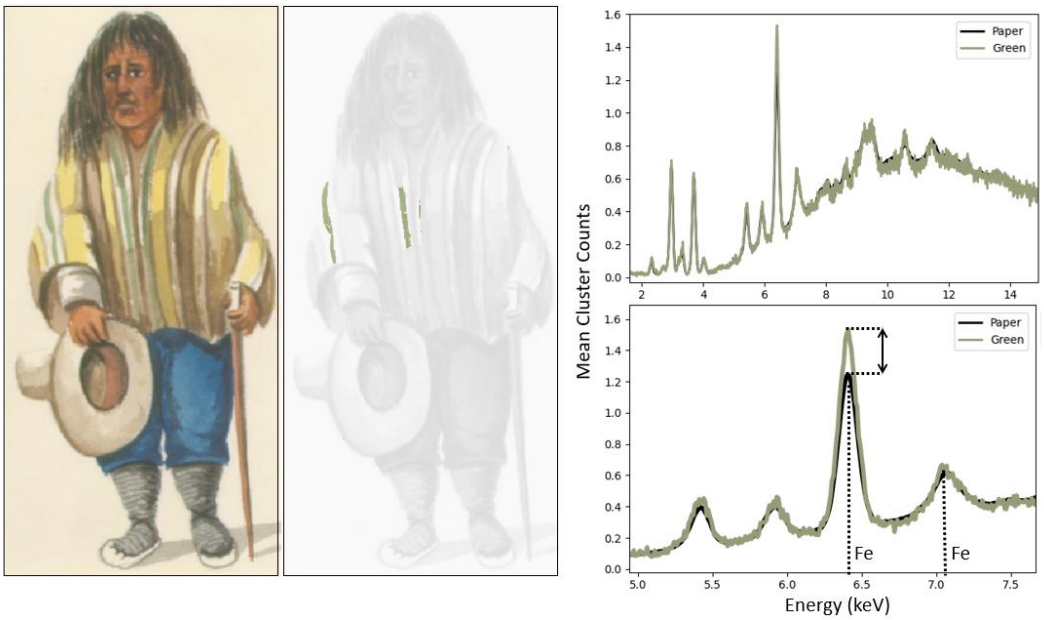

Fig.4.13 Representation of the mean cluster spectrum for the now correctly clustered green shirt pixels highlighted in the LHS cluster map. The higher SNR spectrum in the bottom RHS now illustrates a clear difference in material which was not previously detectable.

This new outcome and the previously covered results clearly demonstrate how beneficial performing multimodal clustering may be in future studies, or even when returning to previously analysed items. For example, in the context of studying P80001-8, it was previously ambiguous as to which materials were used to create the green pigment on the shirt. However, the provision of this higher SNR XRF data allowed for more accurately guided pigment identification to be performed, where for this green area it was only after understanding that a higher presence of iron

existed that evidence could be found to suggest that the green was in fact a combination of the yellow and blue already used within the painting, shown in the following section 4.2.8, where pigment identification is performed for the painting.

4.2.8 Material & Pigment Identification

With the clusters produced from the multimodal clustering, further grouping of some clusters was performed using both the VNIR and XRF spectra, to allow for basic material and pigment identification to be carried out for different regions.

Paper

Paper cannot be analysed fully using VNIR as there are no clear features which are indicative of specific materials. However, using XRF shows a presence of multiple different pigments, giving an indication to the sizing materials used in the preparation of the paper, as can be seen in Fig. 4.14. Some of the elements contained within the paper include Sulphur (S), Potassium (K), Calcium (Ca), Barium (Ba), Chromium (Cr), Manganese (Mn), Iron (Fe), Nickel (Ni), Copper (Cu), Zinc (Zn), and Lead (Pb).

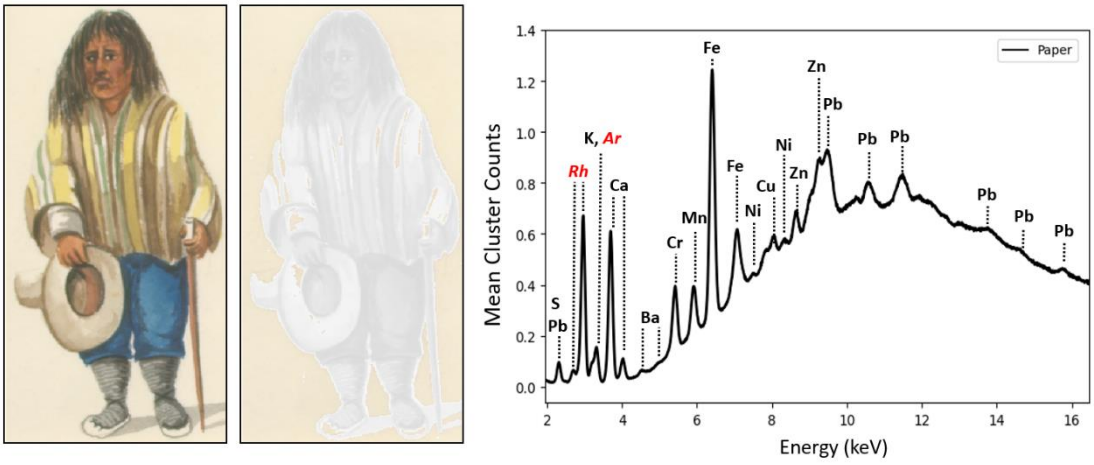

Fig.4.14 RGB and True colour cluster map (left) of the paper in P80001-8, alongside the mean XRF spectrum of the same cluster. Instrumental peaks are highlighted in red.

This many metals being detected could suggest that the paper making process used significant amounts of contaminated water. This may also explain the large concentrations of Chromium which exist only in the paper, as chromium contamination can be common in modern and historical Peruvian water sources (Butler, 2018; Uglietti et al., 2015).

Black & Grey Colours

Over the whole painting there are varying shades of grey or black used to create shadowing effects on the paper or over other pigments, but also to paint the boots of the man. For dark materials, VNIR spectra recorded using PRISMS are not particularly useful, as in the 400-880nm range there tends to be no clear spectral features in most black pigments. Furthermore, as the SNR of most black areas will be low due to its high absorbance, pigment identification is not possible when using Kubelka-Munk. Consequently, the identification of the black pigments must be performed using XRF data, though with multimodal clustering this can be done with higher SNR mean cluster spectra to aid in the identification of both high concentration and trace elements. The VNIR and XRF spectra representing a cluster for the black shoes can be seen in Fig. 4.15. The elemental composition in the black and grey includes XRF energy peaks for lead (Pb) and Sulphur (S) which could suggest the use of black lead sulphide, also known as galena, which would align well with their being many modern mines which produce Galena in Peru today (Gamarra-Urrunaga et al., 2013; Mindat, 2023).

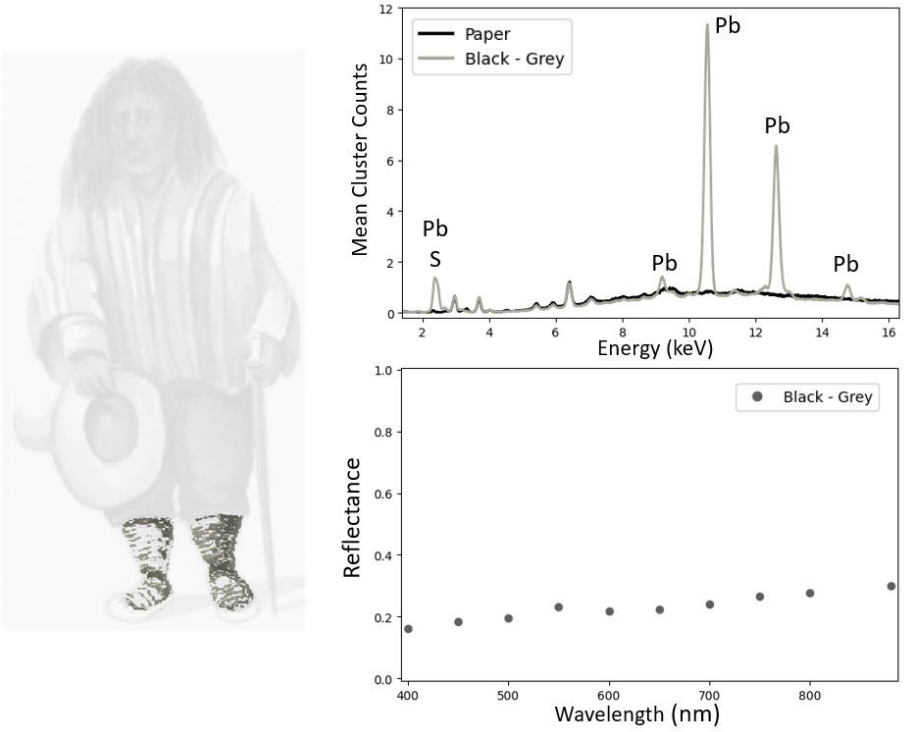

Fig.4.15 True colour cluster map (Left) of the black/grey boots. On the right-hand side are the associated VNIR spectrum, and XRF spectrum which shows clear presence of lead (Pb) and sulphur (S) corresponding with black lead sulphate.

Brown Colours

Multiple areas on the painting could be considered a region which used an actual brown pigment as opposed to adding layers over yellow. These include the cane, the darkest shades towards the bottom of the shirt, but most clearly the hair. An investigation into the mean cluster XRF spectra taken for the brown hair, along with some other regions over the rest of the painting, primarily show increases in the peaks for iron (Fe), Manganese (Mn), lead (Pb), Sulphur (S), and some mercury (Hg) suggesting the use of ochre or potentially umber mixed with some of the previously mentioned black pigment to darken the colour further, with low concentrations of vermillion laid underneath (see Fig. 4.16).

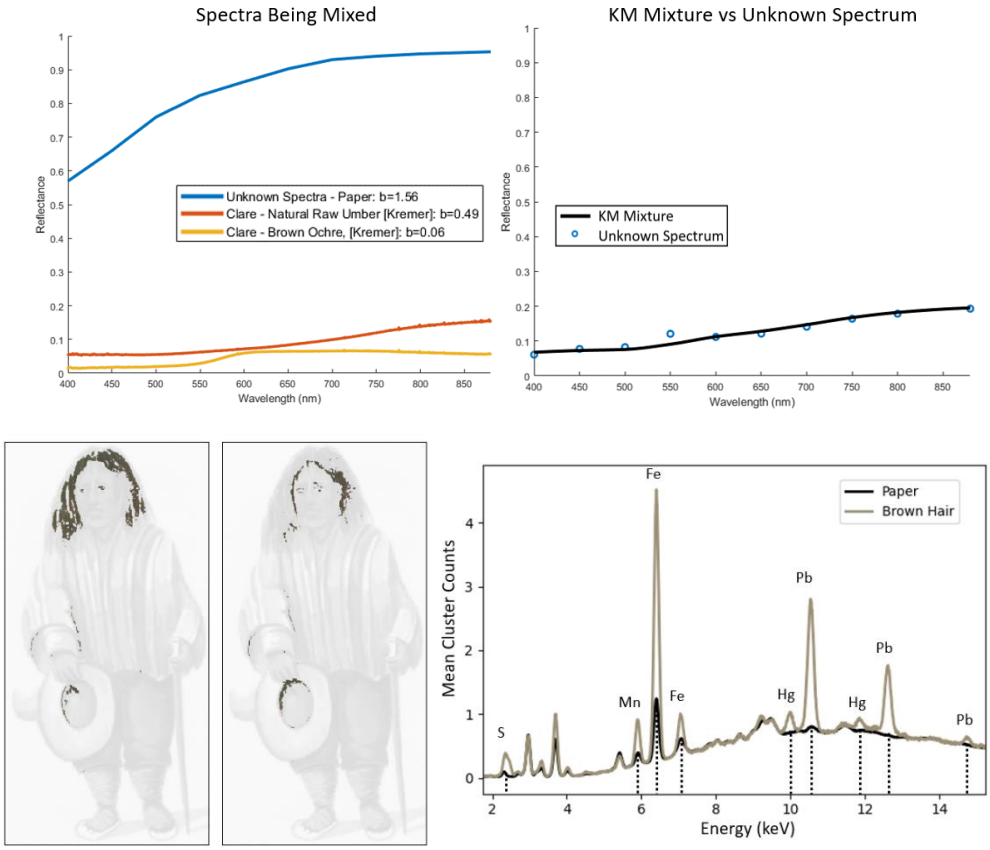

Fig.4.16 Top: Graphs showing the Kubelka Munk mixture for the brown areas. Bottom: images showing the cluster maps (left) for the brown hair and the corresponding XRF spectra comparing the hair and paper (right).

The VNIR spectra provided a dark spectrum which was difficult to fit exactly to any one pigment, however a KM mix of the paper, brown ochre and umber provided a virtual mixture which matched well with most parts of the unknown spectrum, except for the wavelength band at 550nm.

Flesh Tones

As a continuation from the multimodal clustering results, it was shown that an increased presence of mercury (Hg) and sulphur (S) could be detected on the head of the man in P80001-8. After clustering, it becomes clear in fact that the same Hg and S peaks can be detected in all the regions where the man has exposed skin, such as the hands, suggesting that vermillion is used as one of the main pigments in creating flesh tones. When monitoring the VNIR spectra of the different flesh tones, three distinctly different flesh tone clusters can be detected which have changes in VNIR reflectance and XRF spectra. The first is the orange parts of the face and hands which are used as highlights on the brow, nose, and chin etc. The second is the brown-orange parts of the face, which tend to be used as the neutral skin tone of the character. The third is the browner parts which act as shadowing on the face, neck, and hands. The XRF spectra for all three clusters can be seen in Fig. 4.17, and clearly shows that the three regions have the same elemental composition with Hg, S, Pb, Fe, and Mn being detected, suggesting the use of vermillion and ochre, along with the black lead pigment used in various areas of the painting.

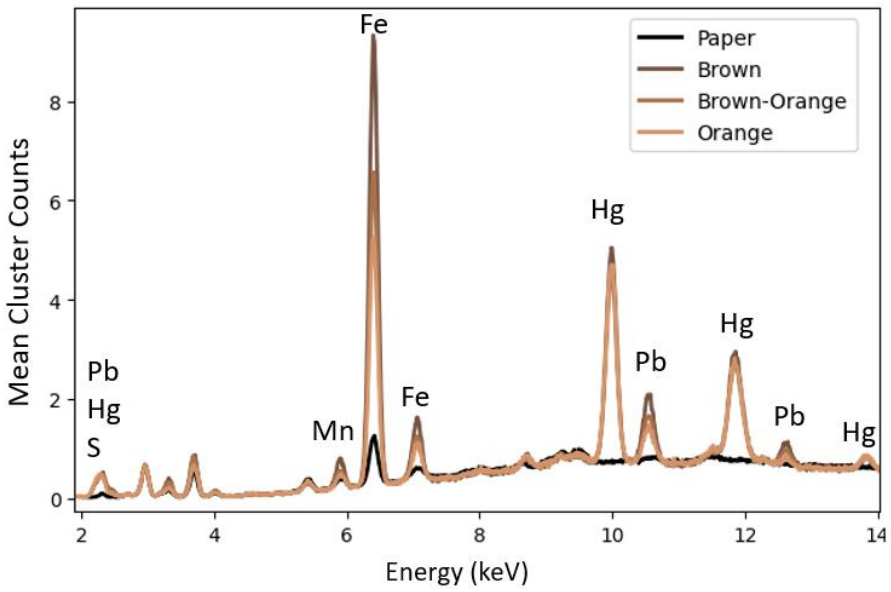

Fig.4.17 XRF spectra showing the elemental composition of different brown areas in comparison to the paper.

To confirm the presence of the different pigments, KM fitting was used to determine the most likely mixtures used for the three clusters, the results of which can be seen in the following sections.

Flesh Tones – Brighter Face

For the brighter face, a combination of vermillion and yellow ochre could be used with the paper to create a virtual KM mixture that had a good fit to the mean cluster VNIR spectrum (see Fig. 4.18). The only deviation for this mixture was at approximately 800nm where the fitted spectrum decreases in reflectance, but the real spectrum does not. This could be due to impurities or perhaps could be a result of the pigment mixture becoming more transparent towards the NIR.

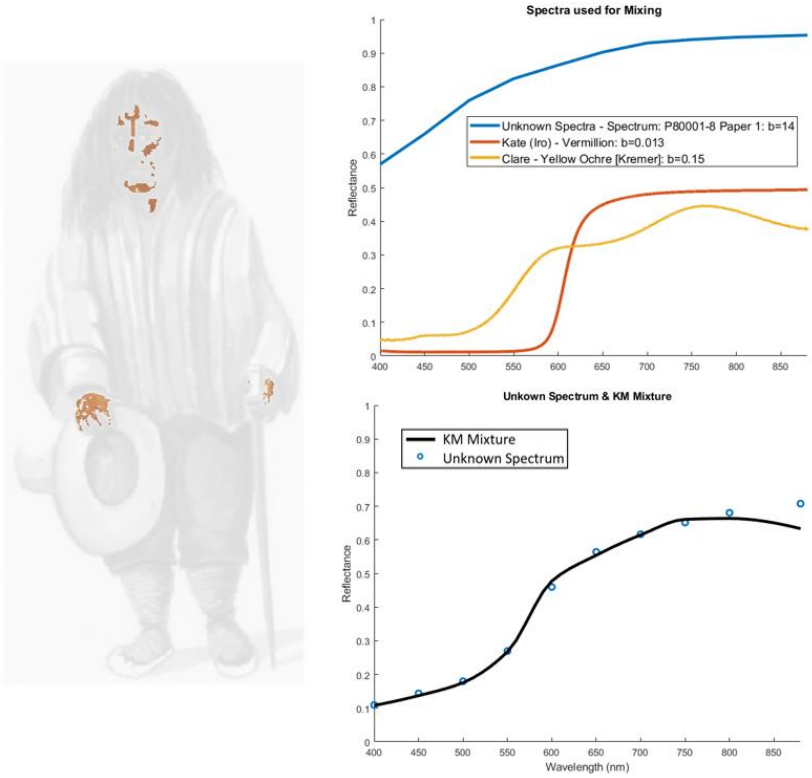

Fig.4.18 True colour cluster map (left) of the brighter parts of the face with the associated mean cluster spectrum and fitted Kubelka-Munk modelled mixture using vermillion, yellow ochre, and paper.

Flesh Tones – Brown-Orange Face

For the more orange/light brown shades of the face a mixture of vermillion and yellow ochre were used again as components in the new mixture, however this time brown ochre was added to the mixture as well. This produced a KM mixture which fitted well with the brown-orange face, however it too deviated from the unknown spectrum at around 750nm. An alternative mixture which used the brown hair spectrum in place of the brown ochre was used and provided a better fit with deviation only at 800nm, suggesting the same mixture was used for shading the face (Fig. 4.19).

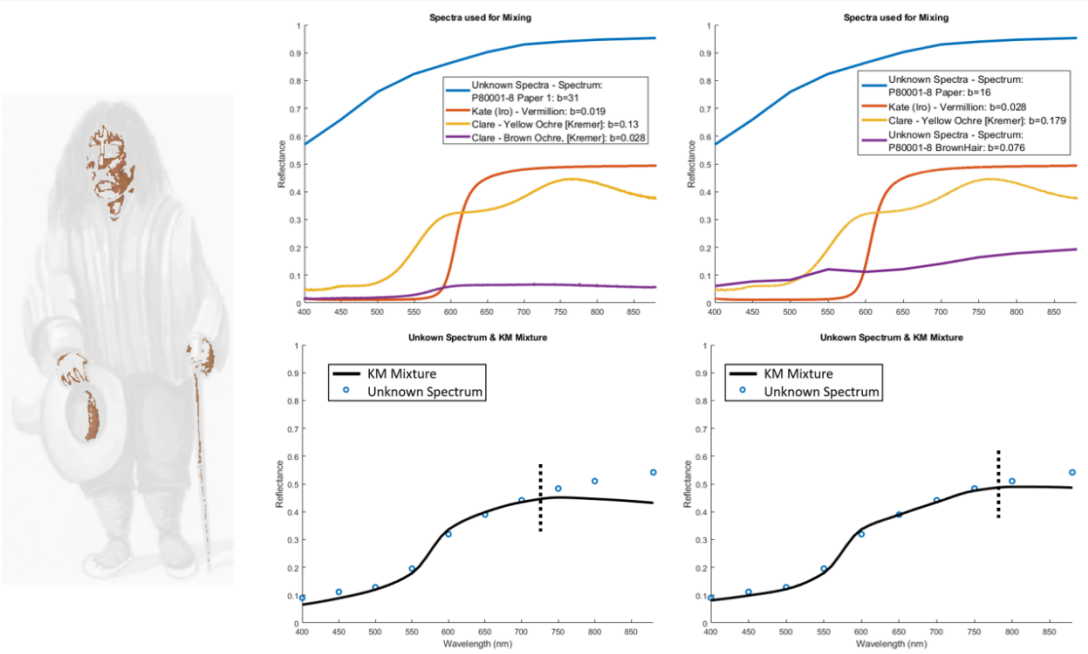

Fig.4.19 True colour cluster map of the orange-brown parts of the face (Left) along with two different Kubelka-Munk fitted mixtures for identifying pigments (Right). The better fitting mixture can be seen on the far-right hand side which used the brown hair as a component in mixing.

Flesh Tones- Brown Face

The darkest parts of the face were shaded with much darker brown than the rest, which could either suggest that more layers of brown paint were used, or a darker mixture was applied directly over the painting. As the brown pigment added to the face was most likely a brown ochre, the same used on the hair, the same mixture used for the orange-brown parts of the face was used again, with the brown hair implemented into the mixture instead of brown ochre (Fig. 4.20). The resulting mixture is a good fit for the darker shades of the face and implies again that the brown is likely applied over the top as an additional layer.

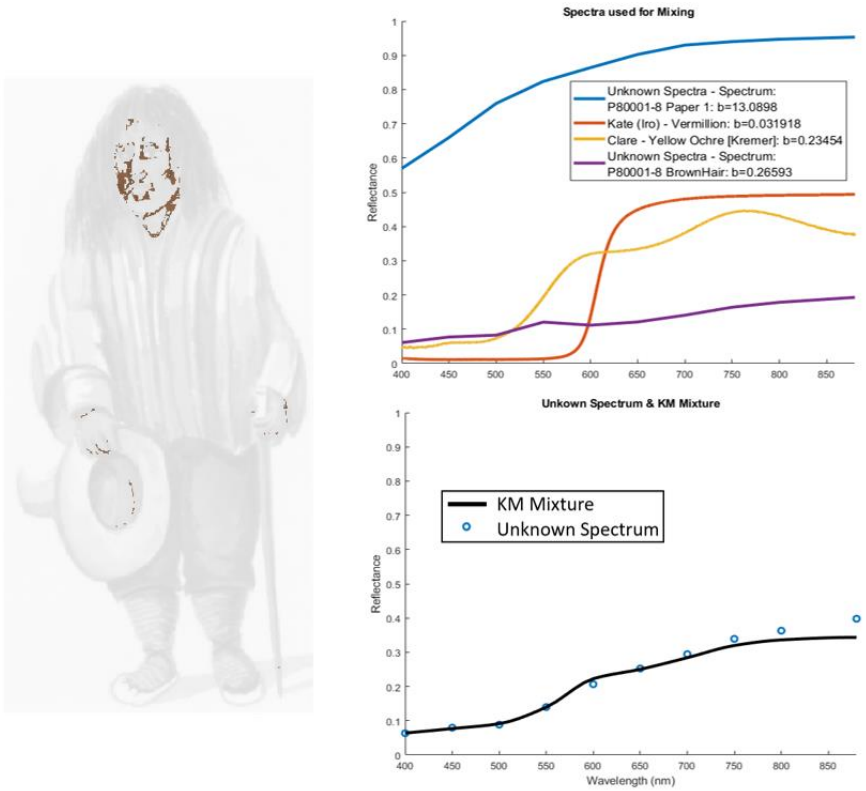

Fig.4.20 True colour cluster map (Left) of the brown parts of the face with the Kubelka-Munk mixture modelling graphs (right) showing that a mixture of paper, vermillion, yellow ochre, and the brown hair, fits well to the unknown spectrum.

Yellow Colours

As illustrated with the XRF clustering, the yellow regions of the painting were noted to consist of only organic materials, as they were grouped together with the paper in the produced cluster maps. Using KM fitting, different mixtures of organic pigments with paper were fitted to the VNIR spectrum of the yellow areas, however for this particular case the best fitting pigment found to follow the yellow the closest was gamboge, which when mixed with the paper would provide a spectrum that closely resembles the painted yellow regions, as shown in Fig. 4.21. It is important to note however that this does not imply that the organic yellow is gamboge, but only that from these results one can clearly see it is likely to be an organic yellow with few spectral features in the VNIR, which at least rules out some other pigments. The only deviation to the KM fit was where wavelength, λ, is greater than 750nm, where the fitted spectrum reduces in reflectance, but the real spectrum does not. However, due to the transparency of many organic yellows towards the NIR, it makes sense that the real reflectance spectrum would increase as the increasing reflectance of the paper would dominate.

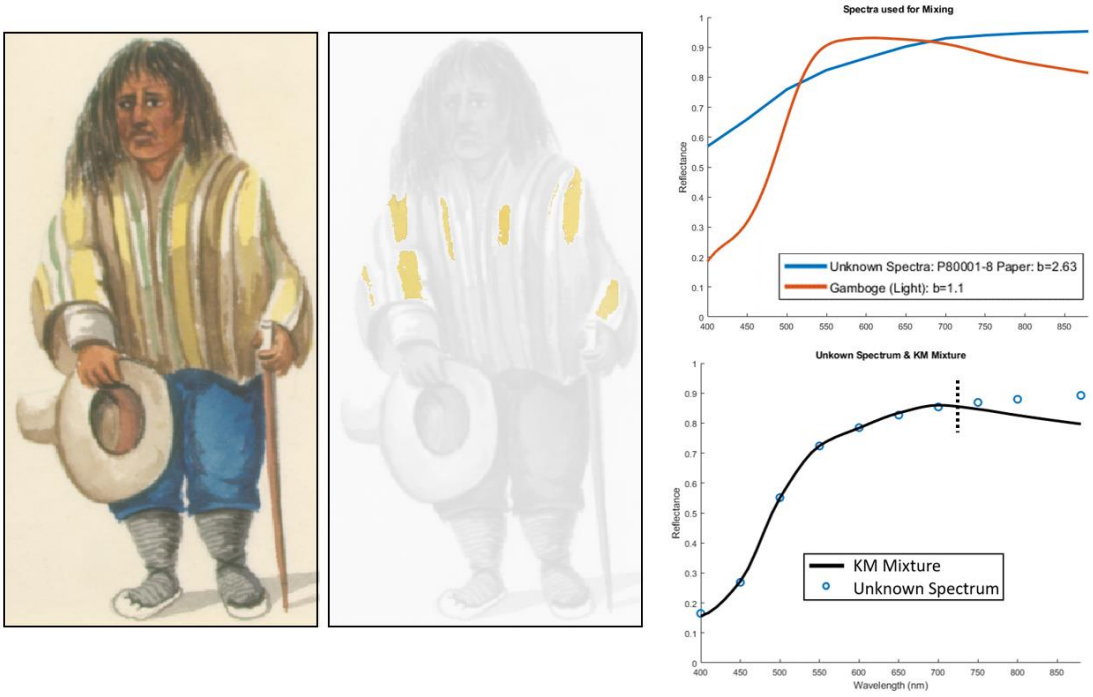

Fig.4.21 Images (left) showing the RGB and true colour cluster map of the yellow regions after multimodal clustering with KM Mixture modelling graphs (Right) showing that a mixture of gamboge and paper fits well to the organic yellow VNIR spectrum.

As it is possible that the yellow used could be any number of organic pigments complementary spectroscopic data could provide insight into the pigment used. Unfortunately, this data was not collected so a definitive identification cannot be performed. This does however show that in future studies, multimodal clustering may provide further benefits if more than just XRF and VNIR are used. For example, if techniques such as Raman mapping or FTIR mapping were also produced

for P80001-8, the separation of specific organic yellow pigments could have been performed during the multimodal clustering stages.

Dark Yellow Colours

As it was suspected that the dark yellow regions seemed to be created by applying varying amounts of a darker shading pigment over the top of the organic yellow, the organic yellow cluster itself was used as one of the reference spectra to investigate how well the unknown spectra would compare with different virtual KM mixtures. When monitoring the XRF spectra for some of the darker yellow clusters on the shirt, increasing concentrations of manganese (Mn) and iron (Fe) can be detected as the shades become darker, suggesting that multiple layers of a brown ochre may have been applied over the top of the organic yellow, furthermore, an increased presence of lead (Pb) also provides evidence for the addition of the black, lead containing, pigment used elsewhere in the painting. KM virtual spectra were produced to investigate the different shaded regions, where it became clear that a mixture of the inorganic yellow with both yellow ochre and brown ochre would create a reasonable fit for different darker yellow shades, as shown in Fig. 4.22. Once again, the KM mixture deviates from the unknown spectrum at λ>750nm, which could be expected if the organic yellow pigment used in the mixture is transparent towards the NIR.

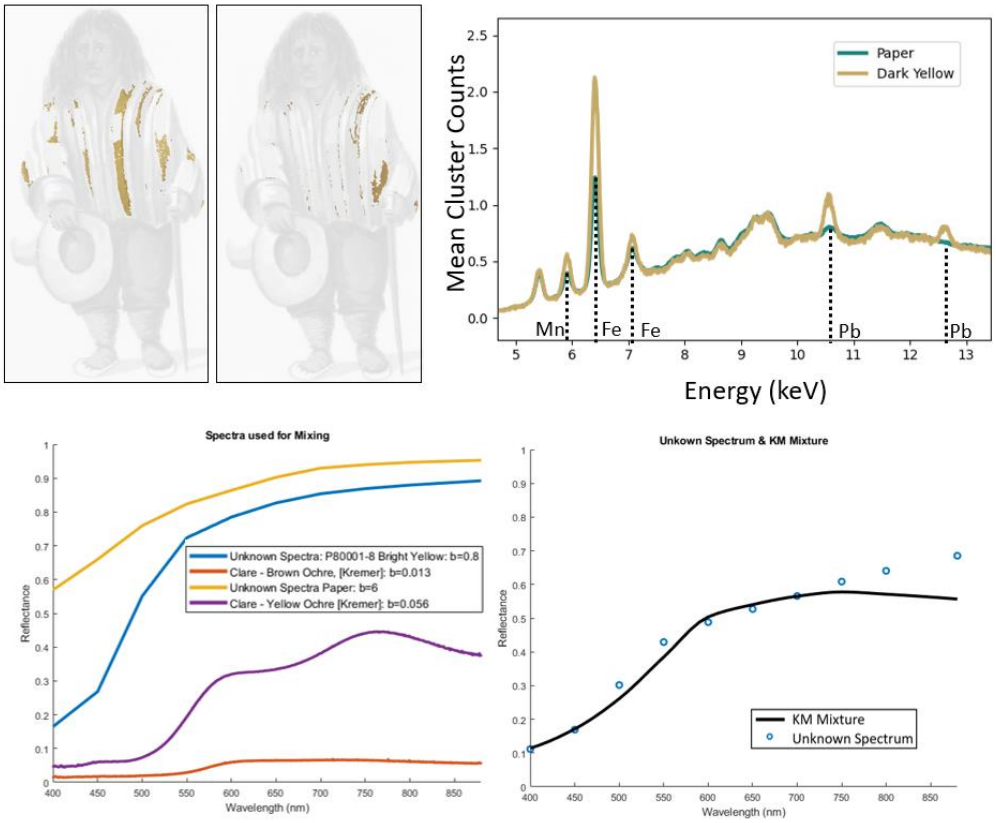

Fig.4.22 Top: True colour cluster maps (left) showing the dark yellow clustered areas of the painting next to the XRF spectrum (right) of a dark yellow area showing the presence of manganese, iron, and lead. Bottom: Graphs showing the Kubelka-Munk mixture use to create a good fitting VNIR virtual spectrum.

Blue Colours

The different blue clusters were joined together into a single group to perform pigment identification, as all the spectra possessed the same broad absorption features. The VNIR spectrum of the differently shaded blue trouser clusters all possessed spectra which fitted reasonably well with a KM mixture of Prussian Blue with paper. The only deviations from this were for some of the darker, low signal to noise areas where it is likely that a thicker application of Prussian blue was applied to paper, sometimes with the addition of another dark colour similar in elemental composition to the black pigment used. This mixture and its darker variation are supported by the mean XRF spectrum which show peaks for iron (Fe) in all the blue regions, with further increases in iron and lead (Pb) in the toward the darker blue areas. The spectra associated with the blue trousers can be seen in 4.23.

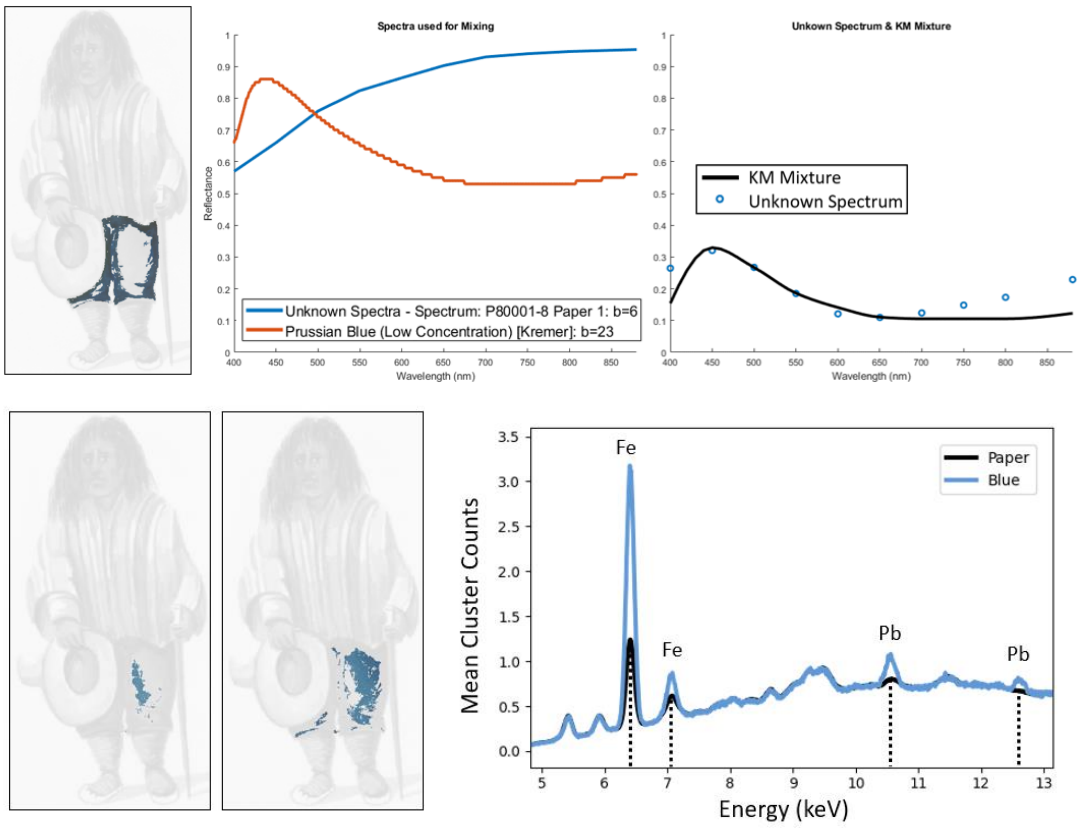

Fig.4.23 True colour cluster maps for the differently shaded blue areas accompanied by the Kubelka-Munk mixture (Top Right) which fitted well to the lowest concentration blue. The corresponding XRF spectrum (bottom right) is also presented which indicates the presence of lead (Pb) and iron (Fe).

Green Colours

Due to the increase in SNR spectrum that was made available by using multimodal clustering, pigment identification for the green region could be performed in a more informative way. Now that the presence of iron was confirmed (as shown in Fig. 4.24), the materials used in creating the green pigment did not have to be inorganic, and therefore KM mixtures with iron-based pigments could be tried with supporting XRF spectra.

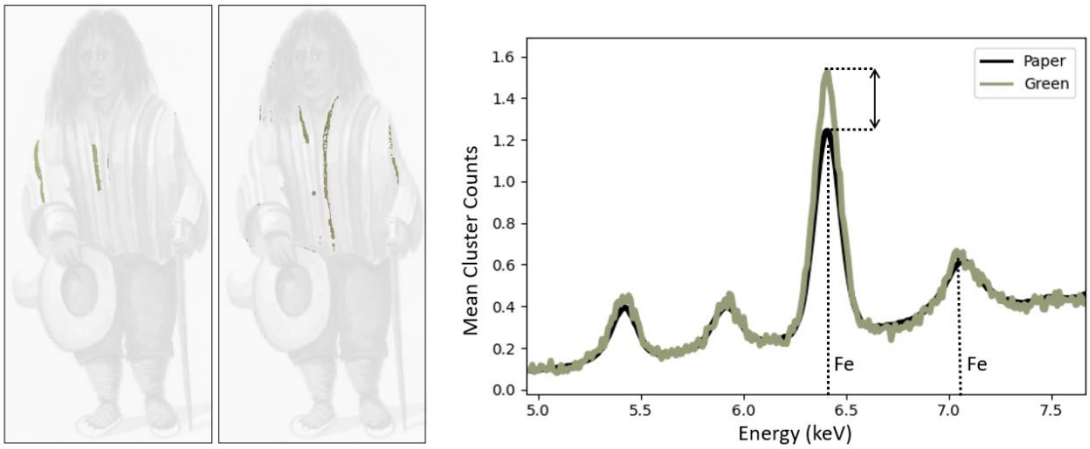

Fig.4.24 Images showing the newly clustered green areas (Left) after multimodal clustering next to the higher SNR green XRF (Right) spectrum which now detects iron (Fe).

As the most likely blue pigment used in the painting was Prussian blue, and the most likely yellow being an organic pigment such as gamboge, a KM mixture using the combination of the two, which can often be commonly referred to as Hooker's green, was used to investigate the unknown green pigment. The KM fit for this can be seen in Fig. 4.25 and shows a closely matching spectrum that begins to deviate at λ>700nm. This may be expected due the applied paint becoming more transparent towards the infrared but can be difficult to model as KM theory begins to produce invalid results for transparent media.

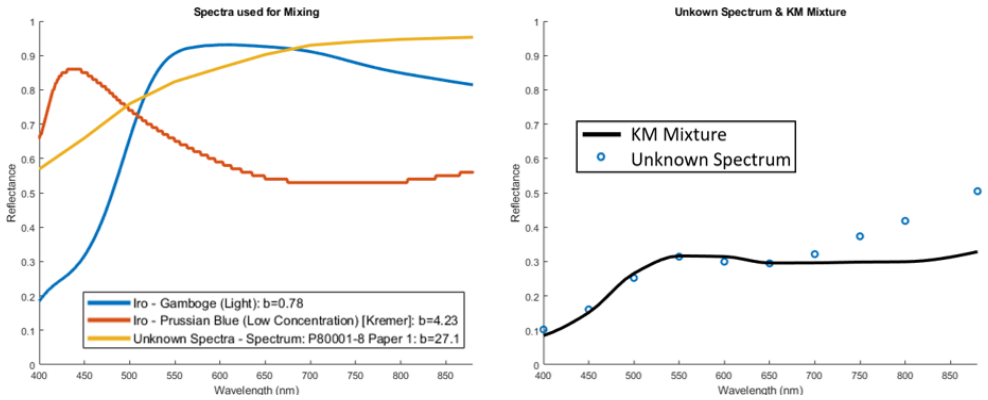

Fig.4.25 A Kubelka-Munk mixture showing that paper, Prussian blue, and gamboge used in as components creates a spectrum which fits well with the unknown green until 700nm, where deviation occurs due to the transparency of the paints.

To alternatively test if the mixture used the same materials present elsewhere on the painting, the two mean cluster spectra for the paper, yellow shirt, and blue trousers, were instead combined using the KM mixture model, with the hope that using the spectra directly from the clustering itself would account for the transparency in the NIR, as the paper substrate used would be identical between the reference and unknown spectra. The result for this KM fit can be seen in Fig. 4.26, which provides convincing evidence that the two colours were indeed used to achieve the green colour seen in the shirt.

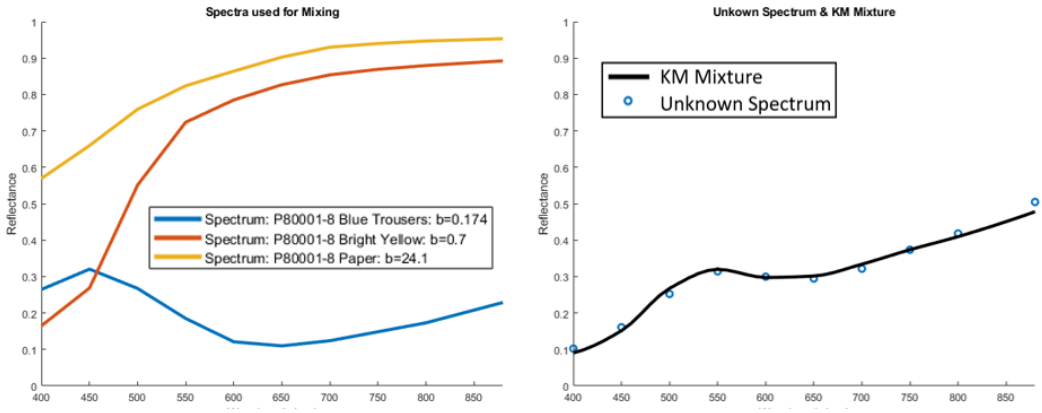

Fig.4.26 A Kubelka-Munk mixture showing that using extracted spectra for the yellow, blue, and paper on P80001-8 can be used to create a well-fitting mixture to the unknown green, suggesting the same pigments are used.

4.2.9 Additional Results & Comments

VNIR Clustering – Detection of Dyes

In addition to the multimodal clustering results seen in the previous subsections, clustering was also performed on the PRISMS spectral imaging data prior to image registration, so that an investigation into the painting could be performed at a more optimal spatial resolution. Performing clustering on this data before registration and using a σ_{VNIR}=1 at the repeated clustering (with the same merging: σ_{VNIR}=2) allowed for the identification of one other pigment not mentioned in the previous sections. On the cheek and lips of the man's face there was an additional spectrum, which had previously been combined into the main cluster for the face, that was suspected to be a red dye. This is a common practise in many paintings and is used to add a 'rosiness' to the face of human characters to produce a more lifelike appearance, so is not surprising to find under closer observation.

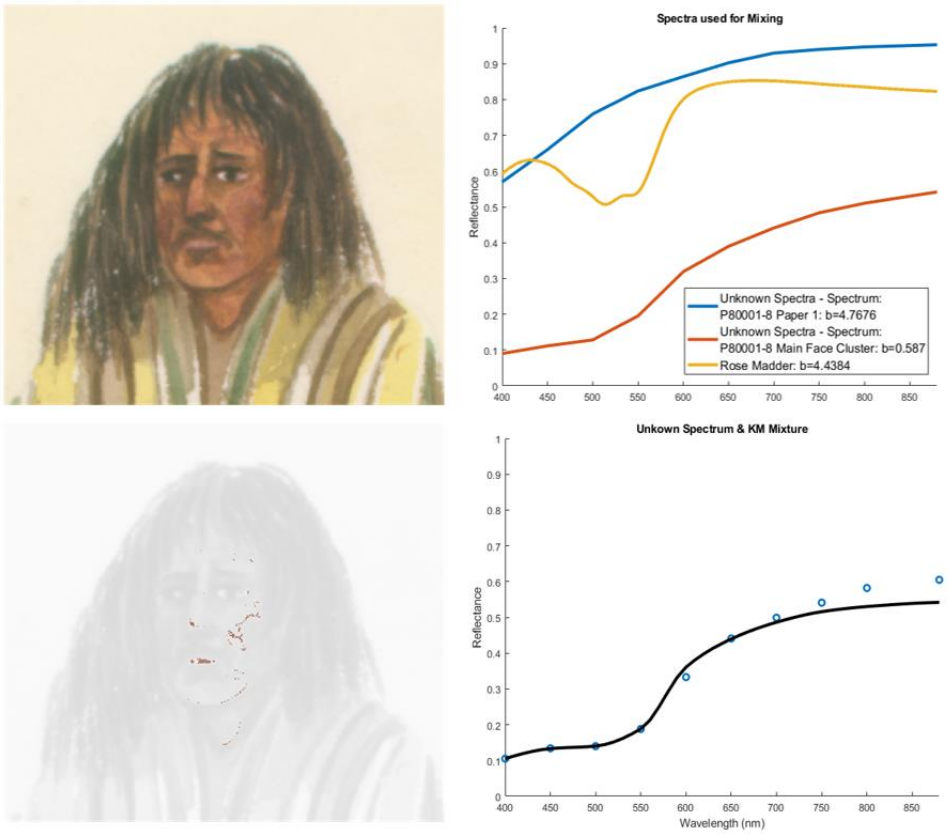

Fig.4.27 RGB Image and representative pixels (Left) which possess a VNIR spectrum which can be estimated to be a mixture of rose madder over the top of the paint on the face when using Kubelka-Munk mixture modelling.

Extracting the spectrum and fitting rose madder as an addition to the main brown/orange cluster for the face produced a virtual spectrum which was a good fit to the unknown red dye spectrum (see Fig. 4.27). This was an interesting result to find, as typically it would be expected that red dyes would only be well distinguished when using a higher spectral resolution, however in this scenario such a distinction can be made by the using the clustering algorithm with PRISMS data possessing a resolution of ≈50nm. This suggests that in some circumstances where fine clustering may be required, such as for the identification of dyes, using smaller standard deviation coefficients could be advantageous, though at the cost of time efficiency as a disadvantage of reducing the standard deviation is that more high precision clusters will be produced.

VNIR Clustering vs Optimal Microscopy

While there were no further additional materials found in the painting in comparison to the registered data, an interesting note to be made for the clustering results of the original resolution data was how well the clustering could be used to clearly isolate spatially small features which were difficult to identify in the sRGB. One good example of this is the distinctive clustering of faint drawings, where the output cluster maps can more easily isolate a finely drawn region such that it possesses sharp spatial features, seemingly with a resolution closer to that of a 20x image recorded using a Keyence VHX-6000 digital microscope (Fig. 4.28).

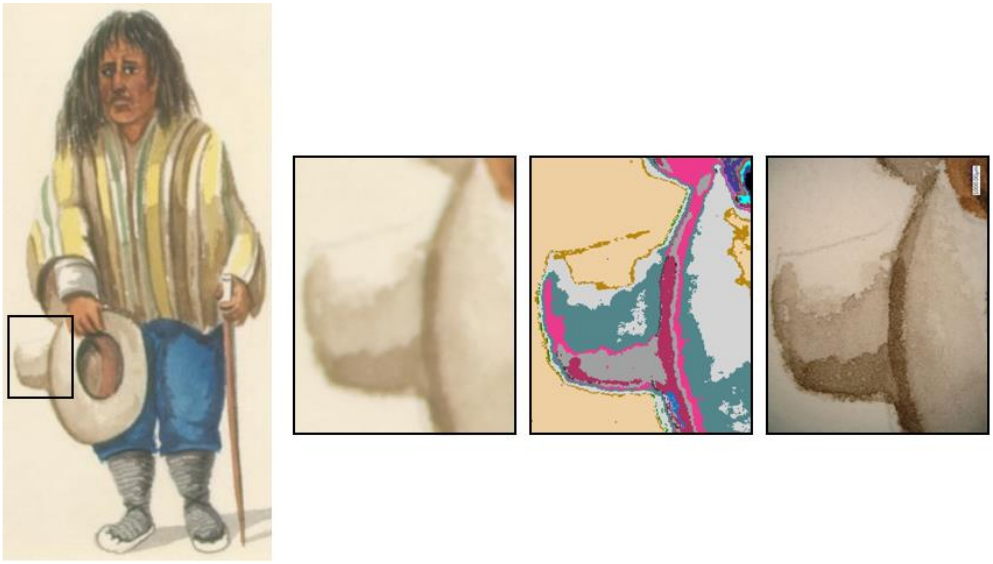

Fig.4.28 Close up comparison of the RGB PRISMS image (Left), with the false colour cluster map (centre), and an image recorded using optical microscopy (right). Notice that the fine pencil underdrawings are more noticeable in the cluster map than the PRISMS RGB.

4.3 Holistic Multimodal Clustering of a Chinese Export Painting

As demonstrated in the previous section of this chapter, multimodal clustering can be used to perform automated complementary clustering between VNIR PRISMS and Bruker M6 XRF spectral imaging data, after which analysis can be performed using a single map with clusters representing unique VNIR spectra, colour and elemental composition. The results from this investigation showed that improvements can be made to the accuracy of the results owing to the inclusion of more than one spectral domain, and that previously hidden information as a result of low SNR or paint layering can be uncovered during multimodal clustering.

Many modern spectroscopic techniques have now developed to the point where they have their own spectral imaging variants, each of which provides spectral data cubes in different modalities which can be brought together to perform a holistic study of the materials in a painting. Therefore, the obvious extension from the previous section would be to investigate whether the algorithm can handle using even more imaging techniques in a holistic clustering approach to eventually allow for the definitive identification of specific materials in artistic works. To investigate such a method, different scientific imaging techniques were used to record data for a Chinese export painting, which possessed multiple artistic materials to create black, pink, red, purple, blue, orange, and green colours. The intention of the different data collection techniques was to perform multimodal clustering across, at minimum, VNIR, XRF and ER-FTIR, so that if the algorithm is used for many techniques in the future, a thorough identification of both organic and inorganic materials could be made. This would provide the information not available previously for the Peruvian export painting, where the conclusive identification of organic yellow pigments was not

possible, while still providing all the necessary data capable of reproducing similar results to those seen before.

Different spectral imaging datasets were recorded for the Chinese export painting on pith paper using VNIR Hyperspectral Imaging, XRF mapping, and ER-FTIR mapping. Complete scans of the painting were performed for the XRF and VNIR to undergo multimodal clustering in the same way as in section 2 of this chapter, allowing for pigment identification for most materials. For the ER-FTIR however, a smaller area was focussed on due to the long collection time.

4.3.1 VNIR Spectral Imaging

It has been well illustrated that clustering VNIR spectral imaging data can be useful for grouping together spectra with similar features, however there are also limitations when performing the clustering on data produced by systems with low spectral resolution, such as PRISMS. Therefore, it may be beneficial to perform clustering on higher spectral resolution data produced by hyperspectral imaging systems so that materials with finer absorption features, such as cobalt blue pigments or organic red dyes, can be properly grouped together. In response, a lab-built line scanning hyperspectral system was instead used to record VNIR spectral imaging data and produced both a colour image and spectral reflectance data cube ranging from ≈400-1000nm with a spectral resolution of 4.5nm, which could both be used as inputs into the clustering and multimodal clustering algorithms (Fig. 4.29).

Fig.4.29 Representation of the RGB and 400-1000nm reflectance spectral imaging data cube collected for the Chinese export painting.

4.3.2 XRF Spectral Imaging & Pre-Processing

XRF mapping was performed using the Bruker CRONO system, where a full scan over the painting was recorded with a spot size of 0.5mm and dwell time of 30ms, with the XRF tube configuration voltage and current set to 50kV and 200µA respectively. The total size of the collected map was 166x131 pixels (83.419mm x 47.273mm) and took a total time of 647.4 seconds to complete. A single slice of the XRF spectral data cube recorded from approximately 10.5keV, and the mean XRF spectrum along with its log plot can be seen in Fig. 4.30 with the different elemental assignments.

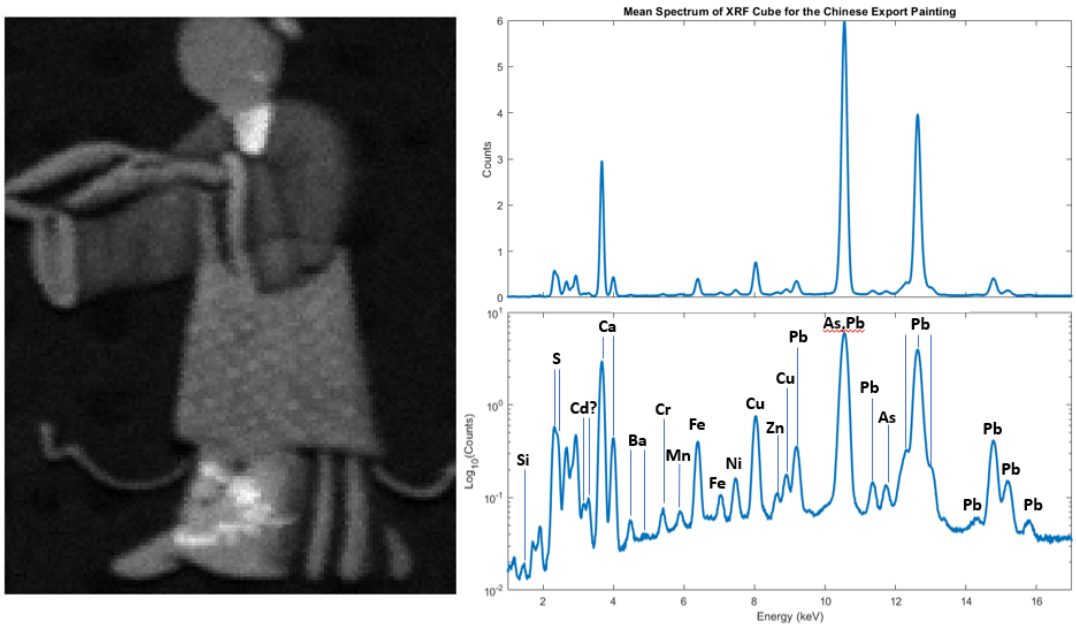

Fig.4.30 Mean XRF image of the Chines export painting (Left) next to the mean XRF spectrum (Top right) and log scaled mean XRF spectrum displaying the identified non-instrumental elements.

In order to perform clustering of the XRF data, the same XRF pre-processing method outlined in the previous section of this chapter was used, resulting in 26 mean XRF channels with contributions from the painting which were combined into a single XRF energy spectral data cube to undergo clustering. The individual channels can be seen Fig. 4.31 after being flipped horizontally to match the direction of other data.

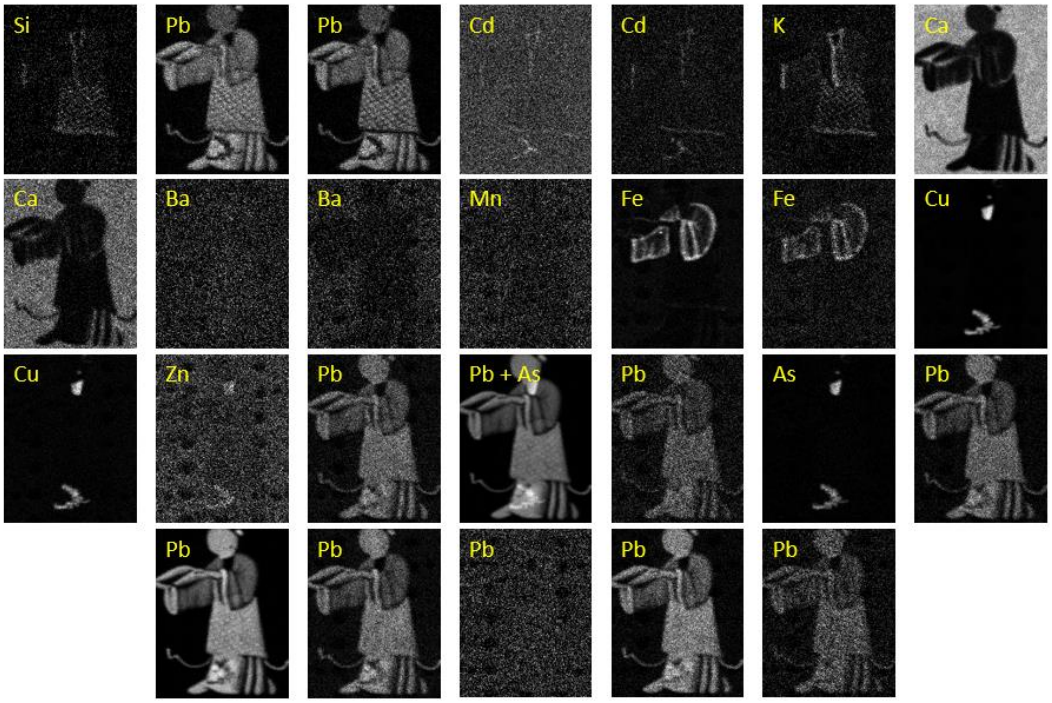

Fig.4.31 Images showing the individual energy channels for the pre-processed 26 channel deep XRF spectral imaging data cube used for clustering.

As can be seen from the individual energy bands created after pre-processing, the data required for clustering in the XRF possessed significant amounts of noise, which when clustered negatively affected results by creating many single pixel "noise" clusters. During clustering, the filtering stages when applied onto the noisy data often meant that much of the painting's cluster map pixels would be removed, so to address these noise problems, a 2x2x1 (2D) spatial median filter was applied to the XRF data cube so that each XRF channel would be free of overwhelming salt and pepper noise. The median filtered energy bands were manually checked to investigate whether any features with fine or sparse structure were lost, and after confirming that this was not the case the resulting median filtered spectral imaging data cube was used as the input into the clustering algorithm.

4.3.3 ER-FTIR Mapping

ER-FTIR mapping is a powerful tool for performing the identification of pigments, and pairs well with VNIR spectral imaging data, as the visible spectra can provide insights to guide the analysis in the IR. If clustering could therefore be performed to automatically group VNIR spectra and FTIR spectra together based on the material composition at any given pixel, data analysis would be able to be performed much more efficiently, while also providing more definitive pigment cluster maps.

For this study, ER-FTIR spectral imaging data was recorded using the Bruker Lumos II to produce a spectral data cube covering a spectral range of ≈600-6000cm^{-1}. The system uses a single objective at 8x magnification with a numerical aperture (NA)=0.6 and spatial resolution of 250µm. This resolution is similar to that of the hyperspectral imaging, making it useful for multimodal clustering between the VNIR and IR as registration would have a minimal effect on the data. As ER-FTIR mapping could be very time consuming and prone to problems during data collection, only a small area of the painting was scanned using the Lumos. An illustration of the recorded ER-FTIR spectral cube can be seen in Fig. 4.32, covering the purple, blue, and green of the robe with some of the flesh tones of the face.

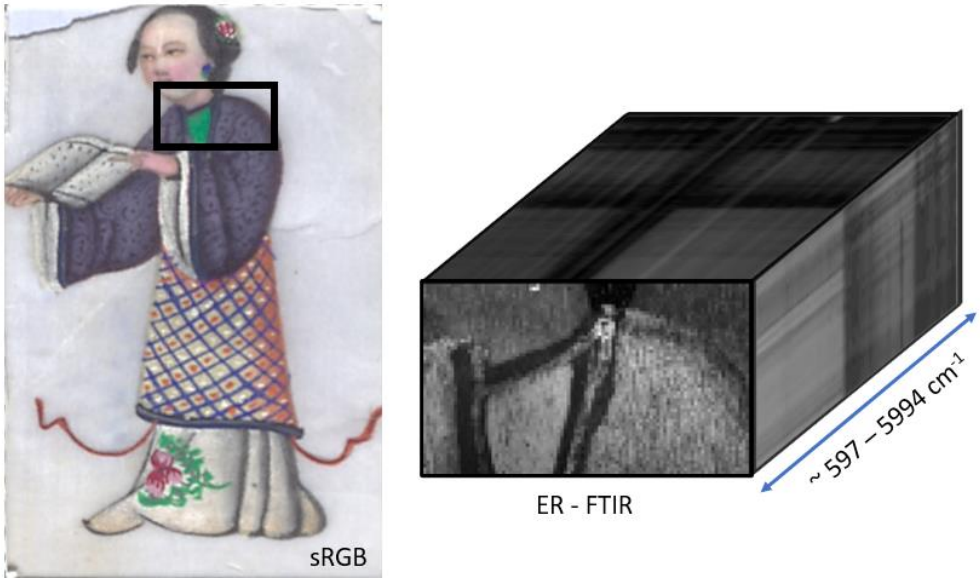

Fig.4.32 Representation of the area of the Chinese export painting scanned using the Bruker Lumos II system in its ER-FTIR mode.

4.3.4 Image Registration

Image registration between different modalities was performed using the same method outlined in the previous section of this chapter where anchor points were selected for common features between spectral images of different techniques before performing an affine transformation (ATF) of the data down to the pixel size of data with the lowest spatial resolution.

VNIR to XRF

As was already demonstrated in section two, image registration can be performed easily between the VNIR and the XRF. Due to the lower resolution of the XRF scanning, VNIR spectral imaging needed to be transformed down to the possess the same spatial pixel size and then cropped to cover the same regions of the export painting. The output images after applying ATF and cropping can be seen in Fig. 4.33.

Fig.4.33 Diagram showing the transformation from the original RGB and VNIR spectral imaging data down to the XRF dimensions and resolution.

VNIR to ER-FTIR

For the transformation between VNIR and ER-FTIR, image registration was performed down to the ER-FTIR with cropping taking place after registration to ensure that both the pixel size and region of interest being analysed between the two techniques was the same. The transformed data can be seen in Fig. 4.34.

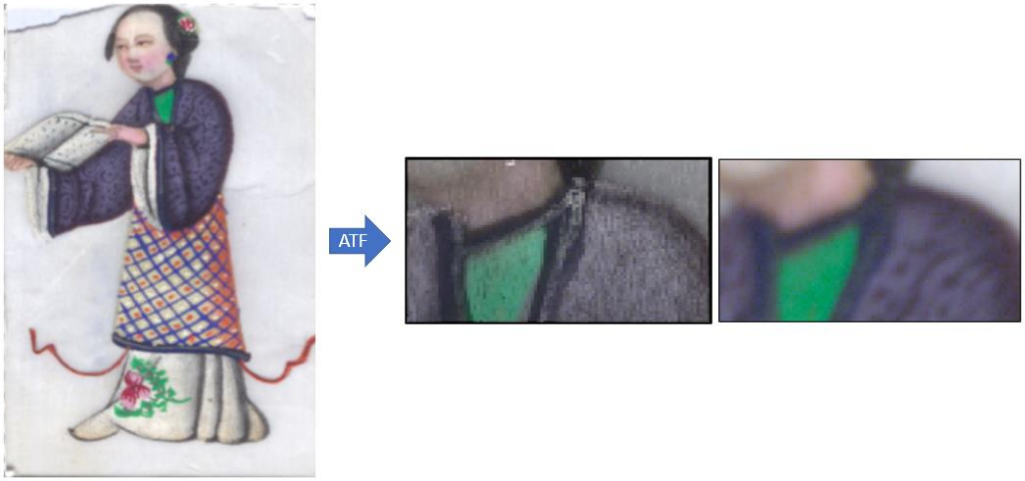

Fig.4.34 Diagram showing the transformation from the Original RGB and vis-BIR spectral imaging data down to the ER-FTIR dimensions, region of interest and resolution.

VNIR-XRF to FTIR

For a more holistic clustering approach, the use of VNIR, XRF, and FTIR together should be able to provide most of the information required to perform definitive pigment identification over a painting. However, in order to carry out the procedure, the data for all three techniques had to first be transformed to be the same pixel size as the XRF and cover the same region of interest as the ER-FTIR.

The biggest difficulty in performing this can be the lack of common features that are visible between FTIR and XRF that could be used as anchor points to have the different techniques to align perfectly for multimodal clustering. The two techniques did not possess many common features, and due to the resolution differences and penetration depth of XRF showing the underpaintings on the pith paper painting, it was difficult to pinpoint solid edges which could be used for registration. However, as VNIR transformations had already been performed between both XRF and ER-FTIR, VNIR spectral imaging could be used as a middle stage to transform between the two data types, allowing for more clear features to be found between the XRF and FTIR by using transformed sRGB images. The different images after conversion to the FTIR region of interest with XRF pixel size can be seen in Fig. 4.35.

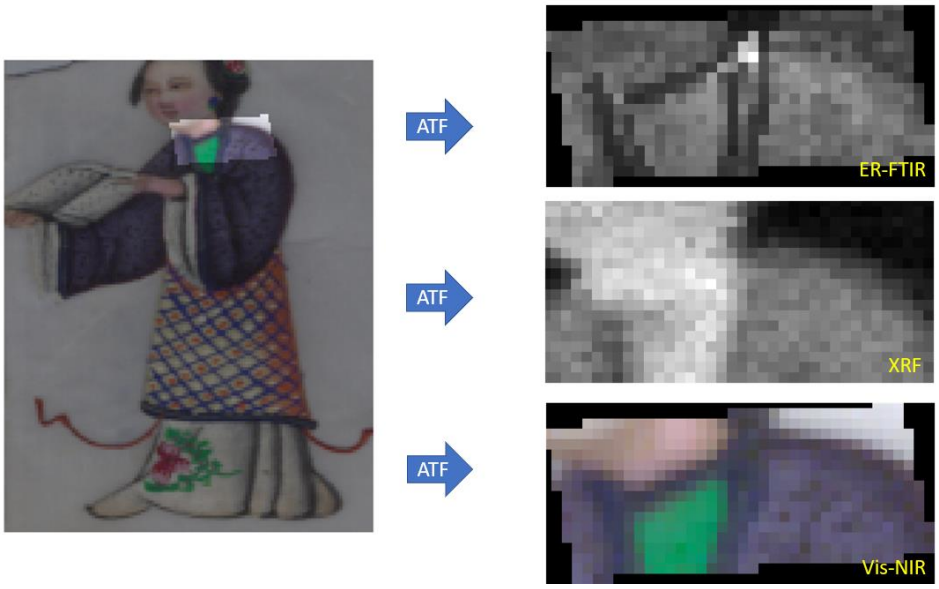

Fig.4.35 Diagram representing the transformations for the three different modalities into the same region of interest, orientation, pixel size and resolution.

From Fig. 4.35 it becomes clear that some of the smaller details of the painting are being lost during the transformation between techniques. This will influence the clustering results, where some information may spread to one pixel from neighbouring pixels and cause subtle spectral differences. This clearly highlights the need to use spectral data of similar resolutions when attempting to acquire all the information, however even with spectral data at this resolution,

multimodal clustering should still be able to perform well. After all of the different transformations had been performed, the individual techniques then underwent single-mode clustering before being joined together for multimodal clustering, the results of which are discussed in the following section.

4.3.5 VNIR & XRF Multimodal Clustering

To gain a full multimodal cluster map encompassing the entire painting, multimodal clustering was first performed over just the VNIR and XRF. For this, clustering using the different techniques individually was first performed so that the cluster maps could be merged for multimodal clustering later.

VNIR Single-Mode Clustering

VNIR single mode clustering was first performed by using the entirety of the VNIR spectral imaging data, however it was found that due to the low SNR of the spectra captured by the system, large variations towards each extreme of the spectral range, caused by low photon counts, commonly caused the clustering results to either incorrectly group spectra by their noise, which may dominate the spectrum, or would produce so many clusters that results would be uninformative. Clustering using the full extent of the VNIR data was therefore not recommended for this study, as it would not be very useful for performing any pigment identification for entire regions within the painting. This made it very clear that in future studies, signal to noise ratio is of paramount importance to ensure that clustering is not affected. It was noted during previous sections that clustering results were sometimes still able to separate dyes from other red pigments at the resolution seen for the PRISMS spectral imaging system. Therefore, for clustering to be performed more accurately for this painting, the spectral imaging data cube was binned into 10 channels, to mimic the same resolution and format of PRISMS data. This 10 band VNIR data was then clustered using parameters of σ_{VNIR}=2.5, and ΔE^*_{ab}=5, producing a total of 108 unique clusters, many of which were redundant but not removed, as filtering the data also resulted in removing important information for an image with spatially small features. A false colour cluster map illustrating the VNIR clustering result can be seen in Fig. 4.36.

XRF Single-Mode Clustering

For the 26 channel XRF spectral imaging data, a total of 51 separate clusters were found when using parameters of σ_{XRF}=2.5, the resulting false colour cluster map can also be seen in Fig. 4.36.

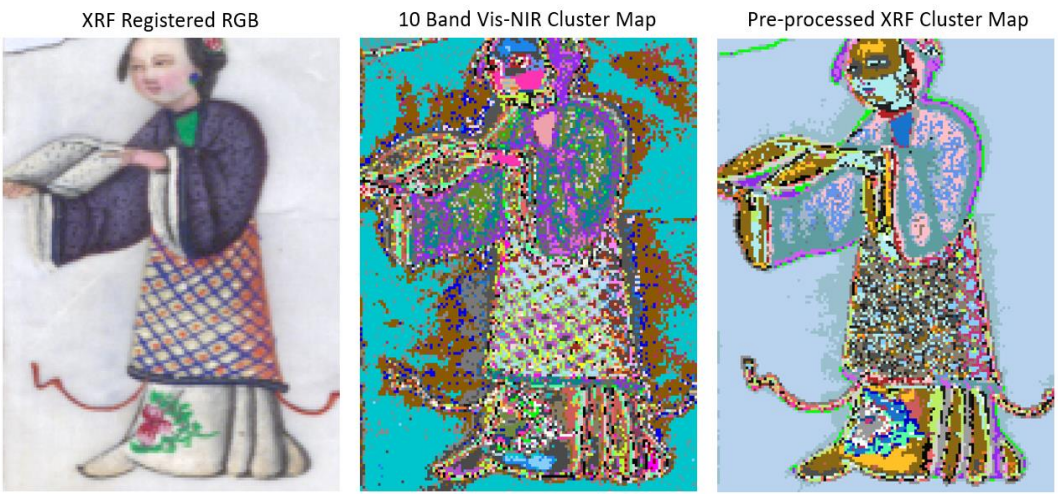

Fig.4.36 Images showing the RGB (Left) alongside the VNIR single mode cluster map (centre) and XRF single mode cluster map (right) both in false colour.

Multimodal Clustering

Using each of the individual techniques, multimodal clustering was performed by using the same parameters as were used for the individual modalities. The results of this provided a final cluster map which possessed 491 clusters in total, which after filtering could be reduced to 60 (see Fig. 4.37). While this number of clusters appears more convenient for performing material identification, having 491 clusters before filtering meant that many of the pixels in the painting that could be useful for pigment ID were removed.

Fig.4.37 Images showing the RGB (Left) alongside the filtered multimodal cluster map representing 60 clusters (centre). On the right-hand side are the discarded "redundant" pixels in true colour.

So, having this many clusters in a small map meant that much of the painting would be filtered out when attempting to remove redundant clusters, but without removing redundant clusters the task of performing material identification across all the techniques was too complex and not time efficient. This meant that either of the approaches would produce a cluster map that is not useful

for performing analysis, however, under closer inspection it was clear that the improved accuracy gained by performing multimodal clustering had also resulted in many of the mean cluster spectra possessing very small standard deviations, meaning that many clusters would not be merged at either the repeated clustering or final merging stages if the same parameters used for each individual technique were used in the multimodal approach. In response, much larger coefficients were given to the multimodal clustering approach to investigate if accurate clusters could be created that would not be removed by the filtering procedure.

Using parameters of σ_{VNIR}=4, and ΔE^*_{ab}=10, and σ_{XRF}=4, resulted in a new cluster map which possessed 51 clusters, very similar to the number produced using the previous parameters. However, the main difference found in this scenario is that more regions now remained after filtering, a good example being the copper-based green decorative leaves on the lower robe, which now appeared as a unique cluster after having been removed during filtering previously (see Fig. 4.38).

Fig.4.38 Images showing the changes between the old and new multimodal cluster map after altering the clustering input parameters. The previously removed copper green cluster can now be seen in white on the new multimodal cluster map and is shown in true colour on the right-hand side.

Of the 51 final clusters, most of the unique pigments found throughout the painting now belonged to a specific cluster, and no clear misclustered data had been produced by doubling the standard deviation and ΔE^*_{ab} coefficient parameters. The result of this test suggests that that even though it would be expected for cluster accuracy to lower when increasing the standard deviation coefficients and colour difference parameters, multimodal clustering instead can offer improved results in some scenarios, likely due to the different spectral domains ensuring accuracy is maintained. However, as some regions, such as the red tassels toward the bottom of the painting, were still filtered out even at this stage, some of the single-mode cluster maps would still have to be used for some regions. Though either way, the new cluster map would still provide useful

information for identifying pigments or revealing other properties of the painting such as paint layering.

4.3.6 VNIR & XRF & ER-FTIR Multimodal Clustering

For the smaller area in with which ER-FTIR had also been used, multimodal clustering could be performed again by first performing single-mode clustering before combining the different techniques into the multimodal approach.

Single-Mode Clustering

For the ER-FTIR registered data, single-mode clustering and grouping was first performed on each individual modality of spectral imaging for the small region on the robe, for which the resulting false colour cluster maps can be seen in Fig. 4.39. Filtering was not used to avoid removing information as the number of pixels was already small.

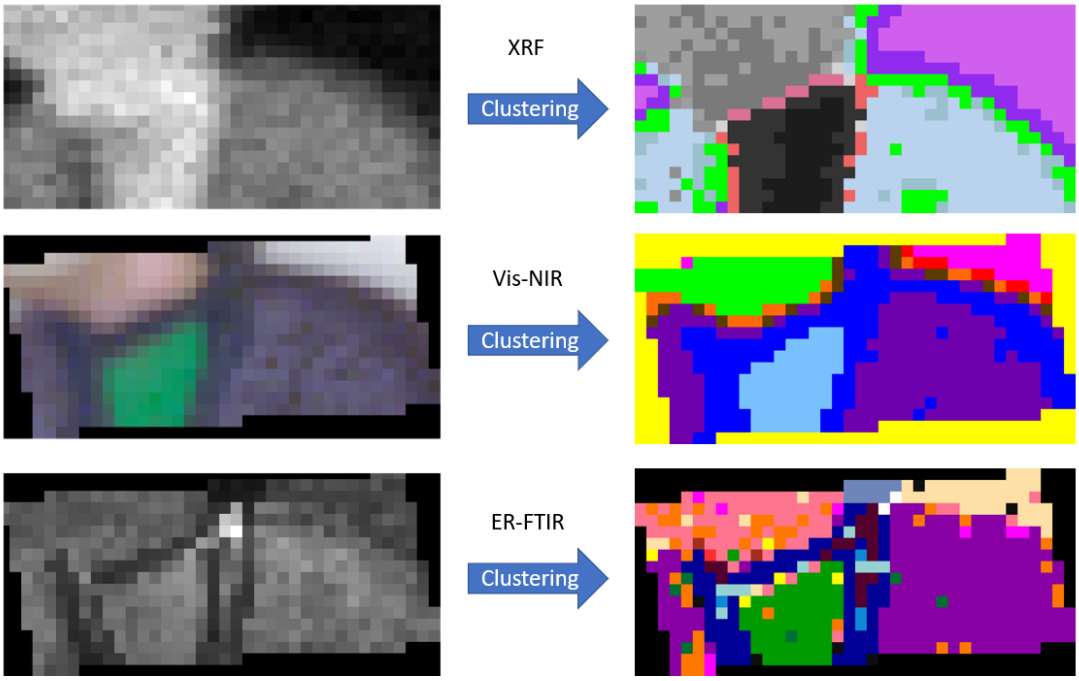

Fig.4.39 Representation of data recorded by the three different spectral imaging techniques before and after performing single mode clustering.

Clustering through each modality produces maps that are similar in structure, however each possesses a slight variation on the different classes. For example, VNIR and FTIR show the features surrounding the green area of the robe very clearly as separate clusters, however XRF does not. In addition, FTIR is more capable of splitting the dark parts of the robe surrounding the

green area apart from the darker outlines, which have been clustered together in the VNIR. The next step after the single mode clustering involved using the multimodal approach outlined in section 2, and used the following free parameter standard deviation and colour difference coefficients to provide clusters, k, as follows:

- ΔE^*_{ab}= 2.5, σ_{VNIR} = 2; k_{VNIR} = 9
- σ_{XRF} = 2; k_{XRF} = 13
- Repeated Clustering $\sigma_{ER\text{-}FTIR}$ = 1.5, Merging $\sigma_{ER\text{-}FTIR}$ = 3; $k_{ER\text{-}FTIR}$ = 17

Many of the different clusters were a result of the interfaces between clustered regions containing pixels which were 'hybrids' of two different areas due to the low spatial resolution, a result of singular pixels possessing signal which would have previously been separated into different pixels before image registration.

Multimodal Clustering

After combining the maps and performing multimodal clustering across all three techniques, a final cluster map was found which possessed complementary information contributed from each individual clustering approach, this holistic multimodal cluster map can be seen in Fig. 4.40.

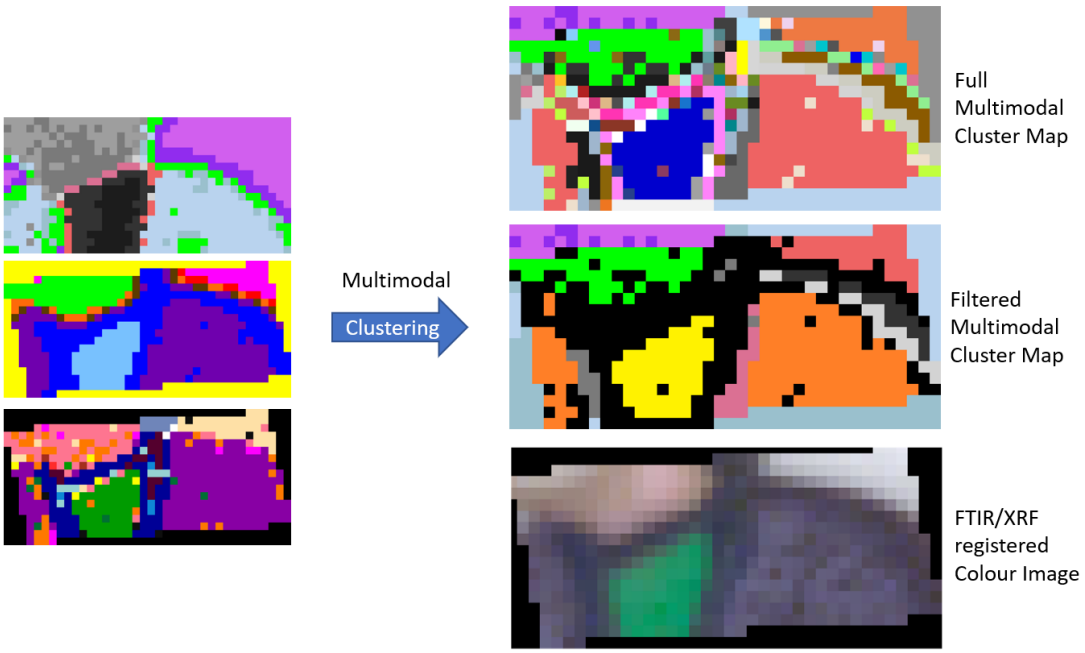

Fig.4.40 Representative before (Left) and after (Right) images of the three techniques being used to generate the full and final multimodal cluster map representing unique VNIR, XRF and ER-FTIR spectra.

After filtering, there were essentially five complete clusters that could be used for material identification, covering the paper, green, purple, skin, and small amount of the dark blue, with each cluster possessing its own unique spectrum in all three techniques. While this small section cannot obviously be used to fully identify every pigment over the entire painting, it can be used for more in depth analysis for these five colours which can still be compared with the VNIR &

XRF multimodal cluster map, thereby building confidence in which pigments or mixtures are likely used over the rest of the painting.

While the multimodal clustering in this example does not provide any further improvements to areas which could have been just as easily be grouped by eye, it is important to reiterate that the demonstration of the 3-technique multimodal clustering on this painting was a practice to test if the algorithm could be expanded to even more modalities. The result of the holistic multimodal clustering producing the same number of clusters as separable by eye shows the validity of the technique and opens the door for applying the approach to new, more complex datasets.

4.3.7 VNIR & XRF & ER-FTIR Multimodal Clustering Results

As the main differences between the clustering performed for the Peruvian export painting in comparison to this Chinese export painting are the inclusion of ER-FTIR and improved VNIR spectral resolution, the results for this section of the chapter mainly focus on what the holistic multimodal clustering offers in addition to that already seen. In this case, this includes the clustering and identification of a red dye used in the woman's face on the Chinese export painting, and confident identification of the green pigment in which all techniques were used.

Red Insect-Based Dye

For the pink colour used on the face of character in the Chinese export painting, it was suspected that a red dye was used. To investigate this, the clustering results showed that the pink areas could be described by a single cluster covering both the cheek of the woman and her hand. While the XRF for this cluster could clearly demonstrate lead (Pb) peaks indicative of lead white (see Fig. 4.41), the additional information provided by the high spectral resolution VNIR showed the presence of red dye in the mean cluster reflectance spectrum, as opposed to the potential implementation of red lead, which would have also been marked by Pb presence in the XRF. Using a KM mixture of cochineal carmine with lead white and paper showed that a spectrum could be produced that fitted very well to the fine absorption lines of the mean cluster spectrum of the pink areas, indicating use of a red insect-based dye (as demonstrated in Fig 4.42).

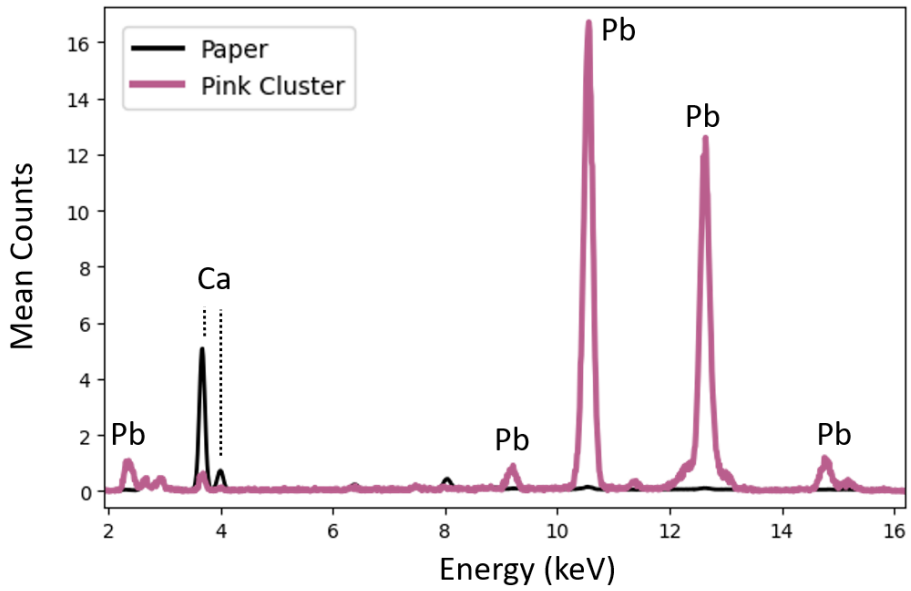

Fig.4.41 XRF spectrum of the pink cluster and paper showing the strong peaks for lead (Pb) indicating the use of lead white.

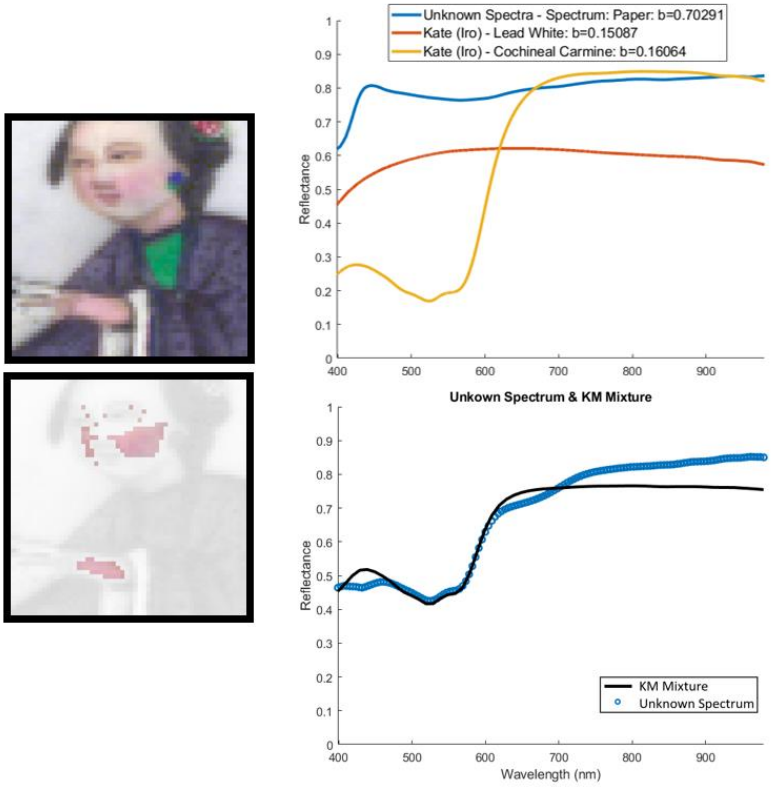

Fig.4.42 Kubelka-Munk fit for the pink cheek and hand of the Chinese export painting, where a mixture of paper, lead white, and cochineal carmine was found to match well with the unknown spectrum.

Emerald Green

With the three technique multimodal clustering results (3 meaning including ER-FTIR), all the spectra necessary to perform a full pigment identification were made available with higher SNR than provided if extracting just a single group of pixels across all techniques. This was best demonstrated in the identification and analysis of the green areas seen in the painting. For example, extracting the mean VNIR spectrum for the upper green area of the robe indicated the potential use of a copper (Cu)-based pigment, due to clear absorption features that matched well with references for malachite or Verdigris (as shown in Fig. 4.43).

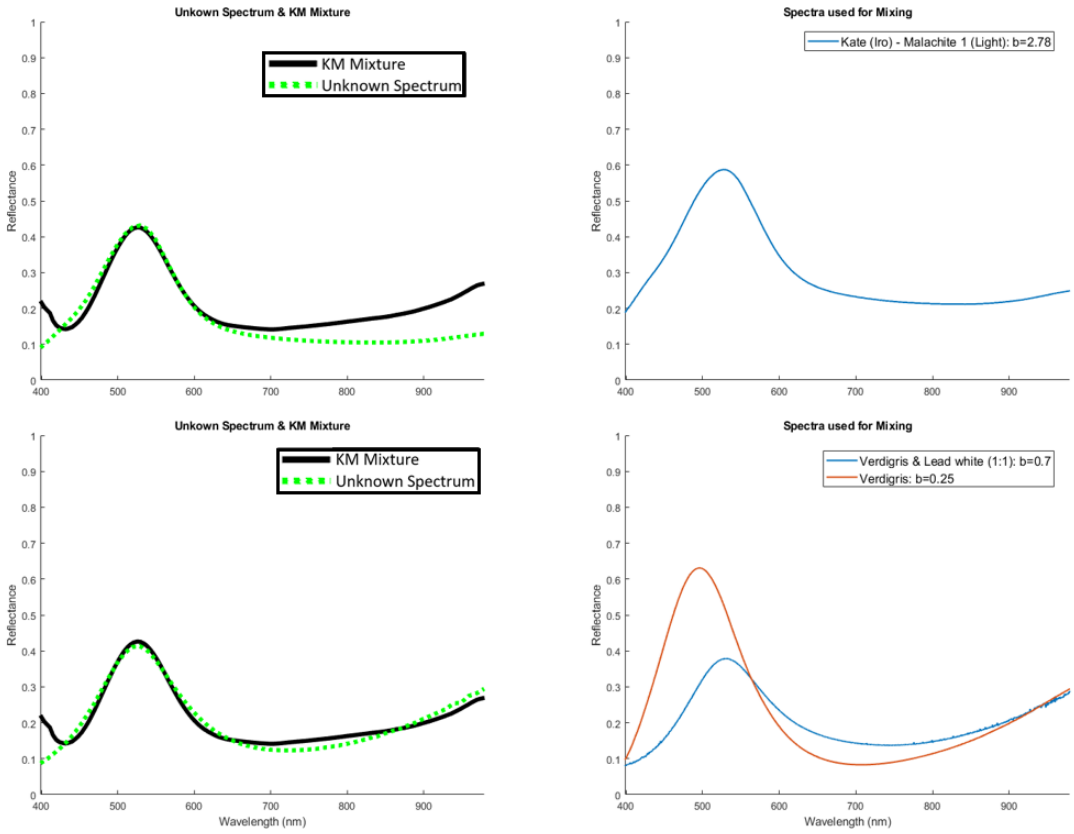

Fig.4.43 Two Kubelka-Munk mixture graphs showing that both Verdigris and malachite VNIR spectra could fit well to the reflectance spectrum of the unknown copper green within the Chinese export painting.

XRF for the same cluster, however, indicated that there was an additional significant presence of arsenic (As) within the area as well, suggesting that the pigment could be a different copper-arsenic based green instead. By using this information to analyse the ER-FTIR mean spectrum for the same cluster in a more informative way, the presence of an ester Cu---C=O-O acetate doublet, found at approximately 2386-2500cm^{-1}, could more easily be identified, which indicated that the green pigment used was Emerald Green. This can be demonstrated in Fig. 4.44, where the VNIR and ER-FTIR cluster spectra can be seen with the region of interest after clustering.

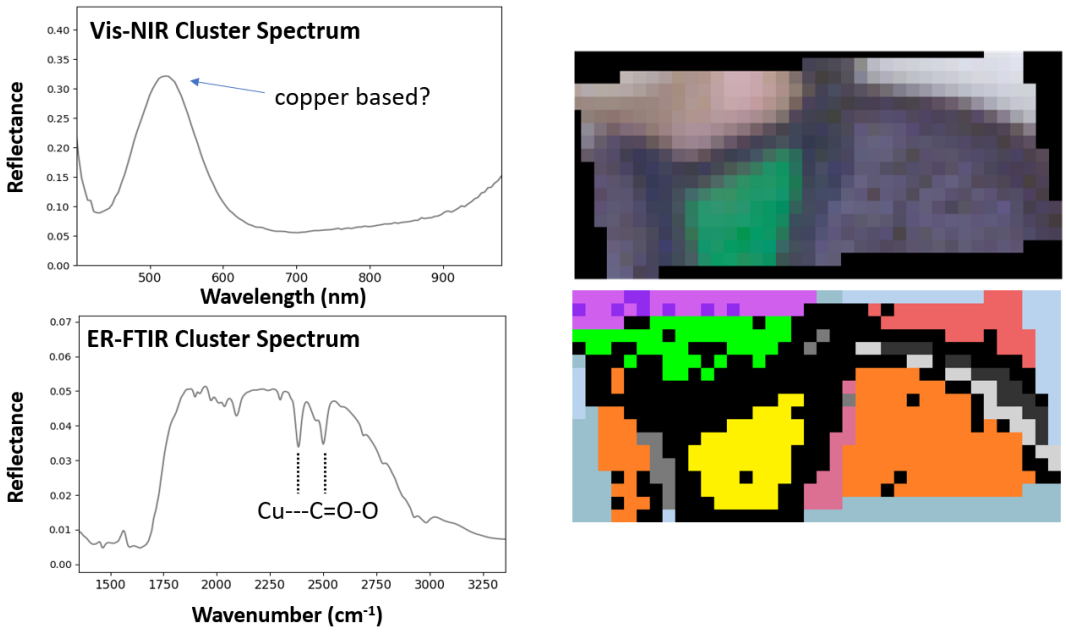

Fig.4.44 Graphs showing the VNIR and ER-FTIR spectrum taken from the yellow multimodal cluster representing the upper green part of the robe. ER-FTIR confirms Emerald Green due to the ester Cu---C=O-O acetate doublet.

Furthermore, in addition to the identification of Emerald Green, using the multimodal clustering results for VNIR and XRF over the entire painting showed additional information which indicated that the upper and lower regions were not the exact same. Under closer inspection of the cluster maps representing the upper and lower green regions, it was found instead that a difference in XRF spectra had separated the two areas due to there being much larger peaks for lead (Pb) in the bottom part of the painting at ≈12.61kev and ≈14.76keV than there were at the top. Furthermore, the presence of copper (Cu) at the top parts of the robe was much higher than at the bottom.

This once again showed that multimodal clustering is able to identify a different layering of paints, where for the bottom part of the painting, the green area was layered over lead white, which was not the case for the upper parts. The associated cluster maps and XRF spectra can be shown to demonstrate this in Fig. 4.44.

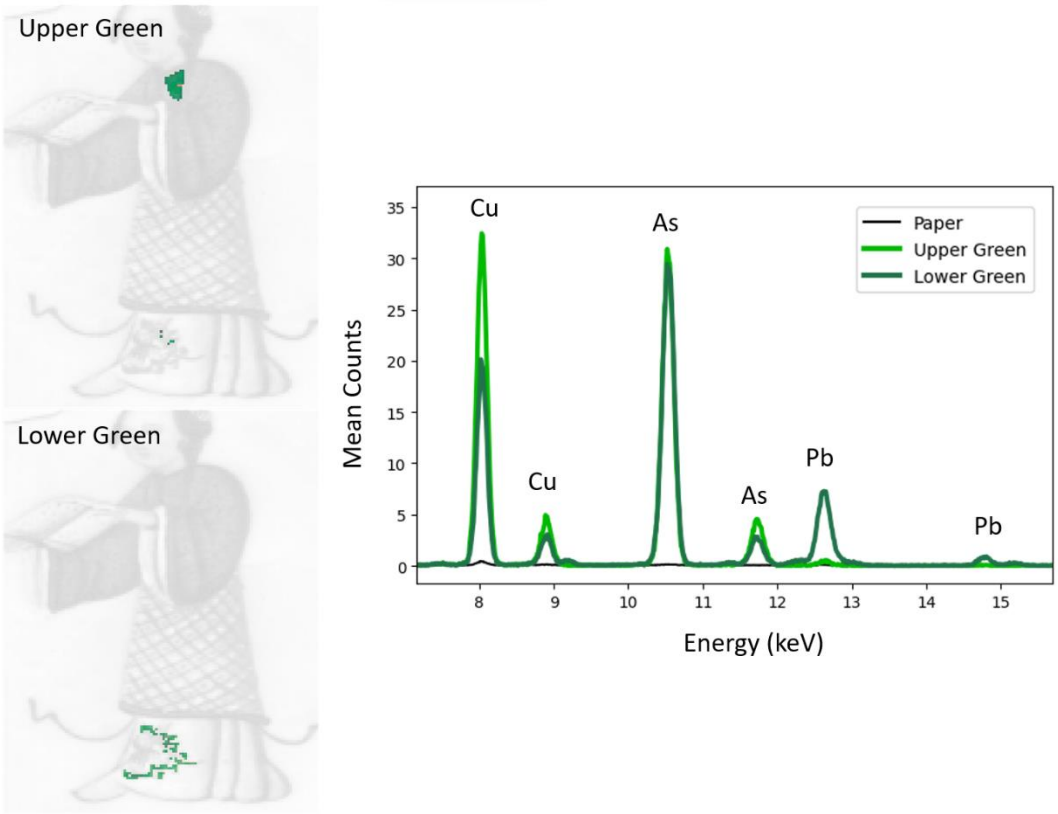

Fig.4.45 XRF spectrum taken from the upper and lower green clusters which demonstrates the presence of

4.4 Conclusions & Further Developments

In this chapter, a novel clustering approach adapted from the new clustering methodology outlined in chapter 2 was developed and tested for performing automated holistic multimodal clustering using CIELAB, VNIR, XRF, and FTIR spectra. The approach developed allows for data fusion and clustering with minimal human interaction with data, where the only hands-on requirements prior to cluster and pigment analysis are to perform image registration across different techniques. This new method was shown to successfully generate cluster maps in which every pixel can be associated with a specific pigment, mixture, or layering of paints, allowing for convenient holistic analysis by analysing only a single image.

During the analysis of a Peruvian export painting, the multimodal clustering method was able to map the distribution and layering of organic and inorganic pigments, allowing for the combination of both VNIR and XRF information into a single map. This cluster map was then used to guide a material identification study and offer insight into the painting techniques used by providing clear illustration of the distribution of pigments and paint layers. Furthermore, the use of multimodal clustering provided mean cluster XRF spectra which much greater signal to noise ratio for areas

which had previously been weak in signal strength at the pixel level, thereby allowing for the more confident identification of thinly applied pigment mixtures.

In a second study on a Chinese export painting, the new method demonstrated its ability to generate cluster maps in more than two spectral domains. By including ER-FTIR mapping data, holistic multimodal clustering could be performed to allow for the complete identification of Emerald green. This allowed for much more accurate pigment mixture maps to be created after clustering. Furthermore, the same method of fusing XRF and VNIR together continued to allow for differences in paint layers to be discovered as well.

Further Developments

The success of this investigation opens new questions into how the clustering of other complementary techniques may be performed in future. In theory, if each individual clustering technique can be performed successfully, the approach of registering and merging multi-technique cluster maps together can be repeated for any type of data. As high spatial resolution imaging can now be performed using a variety of different techniques including short wave infrared (SWIR) reflectance spectroscopy, Fourier transform infrared spectroscopy (FTIR), Raman spectroscopy and more, there may be potential for the clustering approach to be used across many different techniques for many purposes. Furthermore, as the methods used were initially developed for the analysis of collections, the results shown in this chapter open a new avenue for potentially performing holistic multimodal clustering of collections, allowing for much more accurate pigment mixture maps to be created for numerous items at once.

This investigation additionally highlighted the need for higher SNR spectral imaging data to be used in general, so that noise does not dominate clustering results for high spectral resolution VNIR data and so that better separation of thinly applied pigments can be performed. Therefore, in future it would be wise to ensure that high signal to noise data is always recorded as a priority over attempts to collect data in shorter amounts of time, as it is more likely that studies of artistic materials benefit more with post data collection analysis being performed more accurately.

Chapter 5
Automated Classification of Pigment Mixtures in Large Collections

5.1 Introduction

The previous chapters introduced novel approaches to performing more efficient large-scale and holistic studies of spectral imaging data by implementing machine learning-based clustering strategies. With the correct parameterisation, the clustering approach introduced in chapter 2 was shown to successfully map the distribution of pigment mixtures with different concentrations over an entire collection using VNIR spectra, however, when performing the technique, all the spectral imaging data required to be analysed must be used as an input into the algorithm at the same time. In many large-scale studies, spectral imaging data must be captured in situ and over multiple research trips. While one can wait for all the data to be collected for the entire collection before performing analysis, it is more advantageous to perform clustering to map the distribution of different mixtures as early as possible, as the results from a single portion of the data can direct the data collection efforts in subsequent research trips. The implications of this however are that the entire grouping and analysis procedure would have to be reperformed again every time a new dataset is encountered, which could prove cumbersome in large studies and may seem redundant when analysing objects which have already been investigated or may possess the same artistic materials already identified and mapped. Therefore, the ideal scenario for continuing large-scale studies after performing clustering on an original base of data would be to extend the data analysis workflow to allow for the already created pigment groups to be used to automatically identify previously analysed mixtures in new unseen datasets. This way, as databases of spectra are built having used clustering and Kubelka-Munk informed grouping to acquire vast amounts of training data, the same processes would not have to be performed in many different studies on similar datasets.

Machine learning-based classification techniques can be seen as a great candidate for achieving this goal, as they have demonstrated, in many different studies, strong performances when classifying spectral data. In remote sensing studies, a large variety of classification techniques are consistently developed and implemented for the purposes of identifying spectra (Maxwell et al., 2018), and in cultural heritage, many different algorithms can be used for automated classification of pigments such as the Spectral Angle Mapper (SAM) (Chakravarty et al., 2021; Kruse et al., 1993), Support Vector Machine (SVM) (Polak et al., 2017) or convolutional neural networks (CNN) such as the approach covered in Kleynhans et al., (2020). In modern computing and analytical approaches there are now many different supervised machine learning-based models

which could be used for the purposes of performing pigment mixture mapping and identification on unseen data. Additionally, it is not known how well different classification models will perform with the spectral imaging data used throughout this thesis. As such, it is important for evaluations of different techniques to be performed before performing classification on the British Library's maritime southeast Asian illuminated manuscript collection.

5.2 Training & Validation Data

When studying the basic use of pigments for a large collection, classification techniques would be best implemented to identify different mixtures over many manuscripts, regardless of the concentration of the constituent pigments. The most ideal scenario for a large-scale study would therefore involve performing the clustering and grouping only once, to identify unique mixtures for a single or small group of manuscripts that would allow for good coverage of the expected palette. The results of the clustering and grouping could therefore be used as training data for a classification approach which would automatically identify pigment mixtures over the rest of the collection.

This approach outlined in this chapter intended to test whether performing pigment identification on a small group of manuscripts could allow for the eventual automatic identification of the same materials in new objects. This could mean that when performing a large-scale study, new spectral data could be introduced into an already trained classification model to either identify already-seen pigments or reveal new materials which would require further study. Not only would this allow for pigment identification to be performed automatically, but such an approach would also save time during in-situ research as it would indicate where to focus future efforts in performing more in-depth point measurement analysis. Furthermore, as training is performed by using clustering results directly from cultural heritage objects, underlying relationships between the concentration and application of pigment mixtures in training and new unseen data may be uncovered, which would not be the case if pure samples or endmember spectra are used for classification.

While there are many different classification techniques which could be used to achieve our goals, it was unknown which model may perform best for the VNIR reflectance spectral data used throughout this thesis. Therefore, prior to implementing any classification models for large scale pigment identification, different models had to be trained and tested.

Training Using Javanese Manuscripts

The training data set was produced by clustering and grouping spectra for multiple Javanese manuscripts, ensuring to capture of a variety of the palette used throughout the collection. To ensure a sizeable training dataset could be acquired, 90 spectral imaging data cubes taken over

four Javanese manuscripts were clustered separately from the rest of the collection using the methods outlines in chapter 2. Using a colour difference parameter, ΔE^*_{ab}=2.5, and standard deviation coefficient, $\sigma_{Vis-NIR}$= 2, a total of 5598 unique clusters were found after the automated clustering had been performed. After inspection, no misclustered data could be detected except for regions where it may be expected, for example, when pixels for gold and yellow painted areas were mixed due to their similarities in spectrum and colour, or the same for black text or drawings and the black background.

Using a software-driven approach using Guisi (covered more in Chapter 6), clusters were grouped into 14 different unique classes, each with its own unique material or paint mixtures, for which some example cluster maps can be seen in Fig.5.1. An additional small summary table describing the different classes can be seen in Fig.5.2, where for clarity, classes are given negative numbers to show that they are groups. To ensure that an even spread of pixels is used to train different classification models, 20,000 randomly sampled pixel spectra were extracted from each class for training purposes, with the intention of using 70% of the data for training and 30% for validation across different methods, as these values are commonly in use within classification evaluation studies for hyperspectral data such as Adelabu et al., 2015; Arjasakusuma et al., 2020; Gao et al., 2020; Nguyen et al., 2021. Pixel spectra were used as they offered a greater training pool than if the mean cluster spectra were used, furthermore, the natural variation and randomness that comes with using individual pixel spectra may aid in avoiding overfitting, where models are so precisely trained to the training data that they underperform in terms of accuracy on new datasets.

With the training data sorted into the 14 different pigment mixture groups (or classes), the classification could be performed for the entirety of the collection, however as already mentioned, an investigation into the best performing classification model had to be performed.

Fig.5.1 Resulting pigment mixture maps next to the sRGB images of four different maritime Southeast Asian manuscripts at the British Library. Pigment mixture maps were created by grouping clusters created after automated SOM-based clustering.

Fig.5.2 Table showing the 14 cluster groups which were used as classes to train and validate classification techniques.

Class (Label)	Pigment Mixture / Materials	Pixel Count
0	n/a	n/a
-1	None	41714496
-2	Vermillion (+ Red Lead)	3966663
-3	Vermillion + Bone Ash White	519734
-4	Orpiment & Gold	480203
-5	Vermillion + Orpiment & Gold	415401
-6	Indigo	2380518
-7	Indigo + Orpiment	2362906
-8	Indigo + Vermillion	319531
-9	Indigo + Orpiment + Vermillion	402414
-10	Indigo + Orpiment + Vermillion (Darker)	854973
-11	Bone Ash White	960694
-12	Gold	5303130
-13	Vermillion + Black	705343
-14	Background + Black	4266783

5.3 Evaluation of Different Classification Models

To evaluate different classification techniques for the automated identification of pigments throughout large collections, multiple classification models were trained and validated using the previously mentioned control sample of 20,000 spectra per class, extracted from the four Javanese manuscripts. The main validation technique implemented when testing different classification models was to perform accuracy measurements, where correctly classified validation spectra, considered true positives, are compared against the total size of the validation data to provide a score as a percentage, where:

$$\%\ Accuracy = \frac{Correct\ Predictions}{Total\ Population} \times 100$$

Percentage accuracy was chosen as it can be conveniently calculated independently of whichever technique is being used to perform classification, thereby allowing for a fairer comparison between techniques. Furthermore, the percentage accuracy can also be calculated for individual classes to understand where the model may be performing better or worse, which is often necessary as if only a single accuracy value is used to assess the model, the evaluation results can be misleading. This is because when monitoring the accuracy over all classes, it cannot be known if each is being predicted equally well or if one class is being ignored by the model. For example, a percentage accuracy of ≈93% may be achieved in a model with 14 classes if every class has 100% accuracy except one with 0% accuracy.

For performing more detailed studies on classification techniques, useful visual understanding of how a model performs can be attained by representing the distribution of correctly classified spectra in a matrix. This is commonly known as a confusion matrix and is normally built to show the predicted values and actual values of classified samples, thereby illustrating not only class accuracy, but also where the incorrectly classified spectra are being placed. A basic illustration of an example confusion matrix can be seen in Fig. 5.3, which shows a scenario for predicting classes on 300 samples evenly distributed over three classes. In this example, class -2, and -3 have not been modelled well, with a significant portion of the -3 labelled samples being misclassified into the -2 class.

		Predicted Class		
		-1	-2	-3
Actual Class	-1	98	2	0
	-2	1	99	0
	-3	8	29	63

Fig.5.3 A basic example of a confusion matrix which illustrates a different classification performance for three classes. Class -1 performs well, but -2 and -3 underperform.

In the following subsections, different classification models are evaluated after undergoing training and validation on the randomly sampled 280,000 spectra after undergoing the 70:30 split for training and testing respectively. For each method, different parameters were tested with multiple results recorded for each model including accuracy and training time.

Though there are many different classification techniques that could be used on the reflectance data, only a handful were investigated in detail in this study. This was because the small selection selected and evaluated here were able to reach high percentage accuracies that were deemed acceptable. After a broad evaluation, the best performing models were then tested in greater detail, as seen in section 5.3.7, by monitoring accuracies on a larger test data set and analysing both the new accuracies and produced confusion matrices.

5.3.1 Supervised Self-Organising Maps

In previous chapters the unsupervised variation of self-organising maps had already proven to work effectively in clustering VNIR spectral imaging data recorded using PRISMS, suggesting that the supervised version may also perform well. Previously, supervised SOM had been used within the approach taken for the clustering of cave paintings in the Mogao caves, China, in Kogou et al., 2020, and while this study primarily used the created topographical neural map created by unsupervised SOM to perform supervised clustering, SOM can also be used for classification purposes as well.

The main concern regarding the performance of supervised SOM however is in how the SOM is used to classify target data after training. Typically, after training an SOM, each neuron in the network is given only a single codebook vector which is compared to target data for the computation of the best matching unit. Therefore, as not all constituent members for each class are taken into consideration during classification, SOM may struggle to classify spectra with different intensities but similar spectral features without performing pre-processing of the data. At this final classification stage, further pre-processing is not preferred, as any changes in spectral

features or intensity can represent true variation between different mixtures, which ideally would want to be maintained at this stage. Therefore, if the number of neurons in the network is kept equal to the number of classes (e.g. 14 classes), the neural map may struggle to correctly train an accurate representation of the data which can be used for classification. SOM however does present a different way in which classification can be performed; if a map is generated that is much larger than the number of classes, a better representation of the training data can often be created, with multiple neurons possessing slightly different groups of spectra but all corresponding to the same class (see Fig. 5.4), as was demonstrated in (Wong et al., 2019).

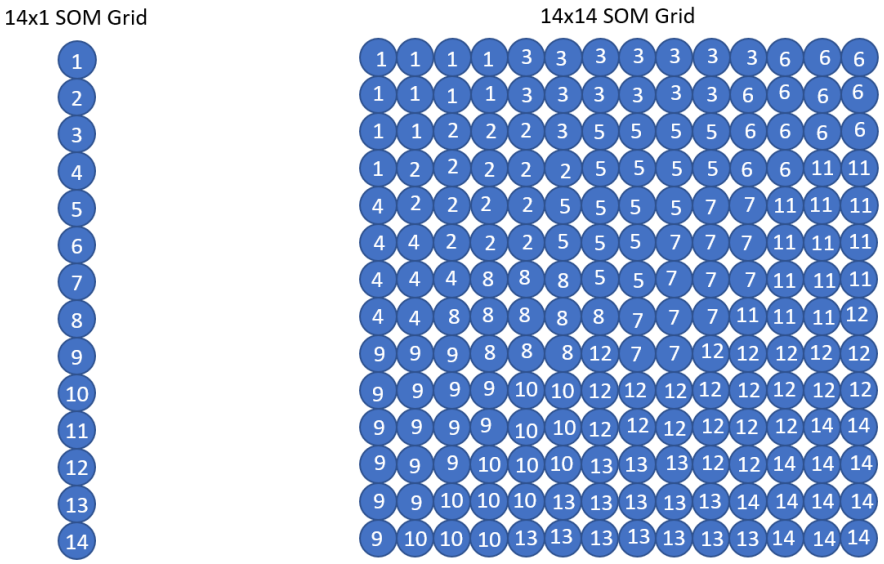

Fig.5.4 Illustrations of two different self-organising maps used for classification. The left-hand side shows the 14-neuron large map, and the right-hand side the 14x14 map.

Furthermore, better SOM accuracy may also be achieved by using different distance metrics. In the implemented "susi" package in python (Riese et al., 2020), four different standard distance metrics can be used to train and classify the neural network: Euclidean, Manhattan, Mahalanobis, and Spectral Angle. To understand which may perform best in a basic test using a 14x1 SOM grid representing the 14 pigment mixture groups, the different distances were compared against each other with increasing numbers of iterations to ensure that convergence was being reached. A summary table showing the results of the training and validation for these distance metric tests can be seen in Fig. 5.5.

Self-Organising Map (Supervised)			
Metric	Iterations	Accuracy (%)	Train Time (s)
Euclidean	1000	19.8	55.19
	5000	23.39	85.11
	10000	31.98	119.81
Mahalanobis	1000	7.15	633.5
	5000	7.35	684.14
	10000	7.44	726.91
Spectral Angle	1000	7.14	118.13
	5000	7.14	143.17
	10000	7.14	180.14
Manhattan	1000	22.55	80.47
	5000	22.65	111.11
	10000	19.83	146.52

Fig.5.5 Table showing the 14x1 supervised SOM accuracy results for different iterations and distance metrics. As highlighted in green, Euclidean distance with 10,000 iterations was found to have the highest accuracy.

While Fig. 5.5 shows that Euclidean distance performs the best out of the different metrics, an accuracy of only 31.98% after 10,000 iterations is far too low for supervised SOM to be relied upon for performing the accurate classification of pigments or mixtures. As it was outperforming all the other distance metrics, Euclidean distance was investigated by itself in greater detail to understand if the poor accuracy could be attributed to the model still not reaching convergence and therefore not being fully trained. To that end, the accuracy for the supervised SOM was compared against an increasing number of iterations (Fig. 5.6).

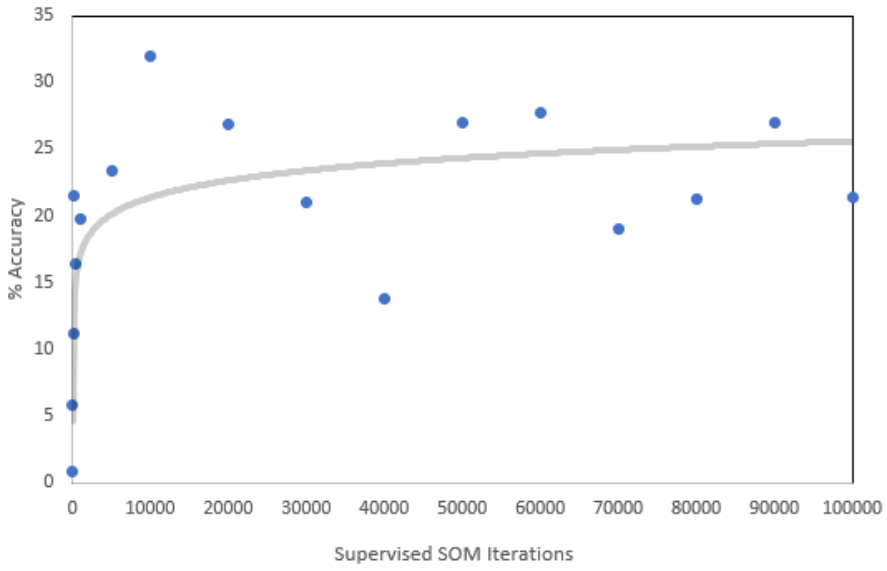

Fig.5.6 Graph showing the change in accuracy with change in iterations for a 14x1 SOM grid.

Fig. 5.6 clearly shows that the supervised SOM does not improve greatly in its percentage accuracy regardless of the number of iterations used past approximately 10,000, as convergence is reached quickly. In fact, the large fluctuations in accuracy after 10,000 iterations indicates that as the number of iterations increases, the SOM does not stay stable either, which indicates it may not be very reliable for supervised classification. As changing the number of iterations does not guarantee a higher accuracy, the results suggest that the only way to improve the accuracy may be to increase the size of the SOM grid so that the increased number of neurons can create a better model of the training data for classification. To investigate the effect of different SOM grids on the accuracy of the model, increasing sizes of SOM grids were created and trained, before being evaluated using the validation data. A graph showing the results of this test can be seen in Fig. 5.7.

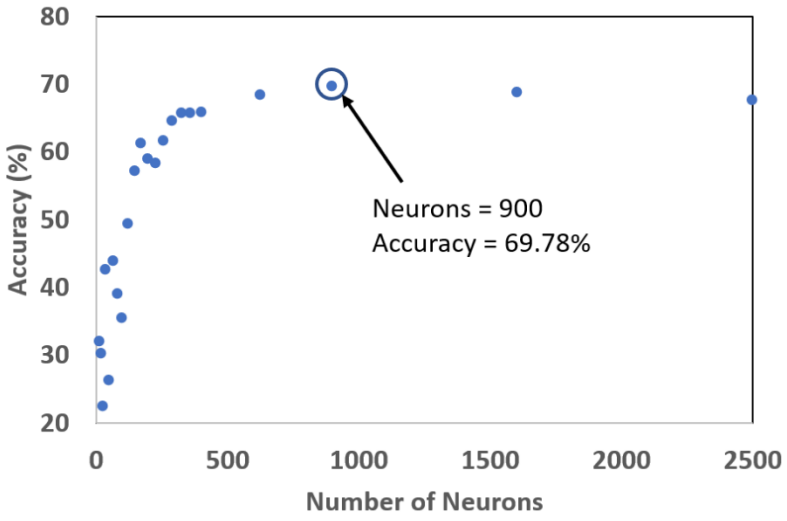

Fig.5.7 Graph showing how increasing the number of neurons increases the percentage accuracy of supervised SOM up to a maximum of 69.78% at 900 clusters.

The results for this final investigation suggest that supervised SOM is simply not a good enough model for performing classification, reaching a maximum of only 69.78% at 900 clusters before only decreasing after reaching a position of convergence, implying overfitting began to occur. This therefore clearly demonstrates the need for alternative approaches which can better model the differences in intensity and spectral shape that can occur due to having different concentrations of pigments in a mixture.

5.3.2 Spectral Angle Mapper (SAM)

One of the most widely used classification models in remote sensing or cultural heritage studies, the spectral angle mapper (SAM) performs classification between "endmember" reference spectra and new unseen spectral datapoints by calculating the angle between the two in feature space. Spectra can be assigned to references by calculating similarities using this spectral angle, where a smaller angle represents more similar data and the larger the opposite. It primarily functions under the assumption that a single reference spectrum, often selected by a user, can be used to accurately describe a single type of material over an entire spectral imaging data cube, a design which aligns well with the classification approach required for the large manuscript collection.

For the application of SAM in the classification of VNIR PRISMS spectral imaging data, the mean cluster spectra were therefore used as 14 representative "endmembers" with which to classify the unseen data. This meant that the resulting classification model would be similar in nature to the implementation of the 14-neuron supervised SOM implemented in 5.3.1, where SOM instead uses codebook vectors learned from the training data to perform classification. While SOM however was only improved by increasing the grid, effectively moving away from this same ideology of classification, SAM does not require significant adaptation of any neural network to increase accuracy. Therefore, only parameter being altered in this investigation was instead the "threshold", which is the maximum value that an angle can have for an assignment of a datapoint to a reference to take place. This threshold value was altered from 0-1 in steps of 0.1 to understand how the percentage accuracy would change with different parameter setups. The resulting table can be seen in Fig. 5.8.

Threshold	Accuracy (%)	Time Elapsed (s)
0	0	0.222031355
0.1	45.2624436	0.223040819
0.2	56.33043251	0.227131605
0.3	56.89591543	0.221031427
0.4	56.96258289	0.224735022
0.5	56.96972583	0.226644278
0.6	56.96972583	0.224033833
0.7	56.96972583	0.237620354
0.8	56.96972583	0.228037596
0.9	56.96972583	0.226085901
1	56.96972583	0.224032402

Fig.5.8 Table showing the performance of SAM under different thresholds. Note that by a threshold value of 0.5, the percentage accuracy had reached convergence, as is highlighted in green.

These results clearly show that while the SAM becomes more accurate as the threshold is increased from 0, it reaches a position of complete convergence by 0.5, with a maximum accuracy of only ≈56.97%. This indicates a worse performance than SOM, which was already deemed not accurate enough for performing pigment classification in this way. As a result, it would be unwise to use the spectral angle mapper for the classification of pigment mixtures in new unseen data, instead implying that a completely new approach must be taken.

5.3.3 K-Nearest Neighbours (k-NN) & Nearest Centroids (NC)

Two potentially useful classification methods that may perform better than the supervised SOM is the k-Nearest Neighbours (k-NN) and nearest centroids algorithms. As the two techniques measure similarities between target and trained data by monitoring the proximity of individual points or centroids in feature space, they should not fall into the trap of trying to fit unknown spectra by using only a single vector in comparison (which in SOM is performed using the BMU). The two approaches may therefore be useful for classifying spectral data that has significant differences in intensity or shape but might consist of the same materials in different concentrations, however, to be sure the two techniques had to be evaluated.

While NC can be used without any optimisation, k-NN requires parameterisation by selecting the number of neighbours. Therefore, the optimal number of neighbours first had to be determined for the k-NN approach. To determine this, the model was trained and validated multiple times with varying numbers of neighbours in Euclidean space, to understand which number of neighbours would produce the highest accuracy. A graph showing the results of this investigation can be seen in Fig. 5.9.

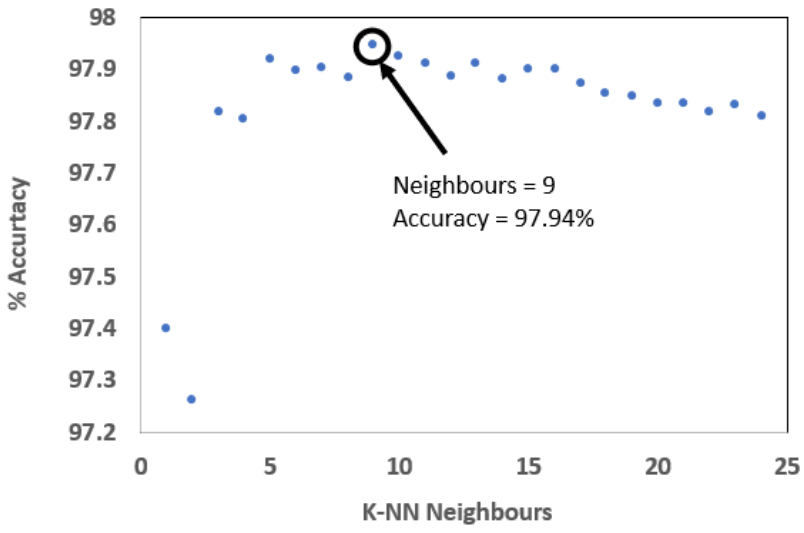

Fig.5.9 Graph showing how the number of k-Nearest Neighbours affects the percentage accuracy of k-NN classification.

The optimal number of neighbours to use the K-NN was 9, and so was kept constant for the investigation between both k-NN and k-NC. The two algorithms were tested using the training and validation data in its 70:30 split, and like with testing the supervised SOM, different distance metrics were also tried for these two approaches, monitoring the accuracy and training time along the way. The associated results table can be seen in Fig. 5.10.

	K-NN (9 Neighbours)		NC	
Distance Metric	Accuracy (%)	Train Time (s)	Accuracy (%)	Train Time (s)
Euclidean	97.94	0.24	75.29	0.026
Minkowski	97.94	0.27	75.29	0.029
Chebyshev	97.59	0.25	77.15	0.024
Correlation	82.78	0.016	68.93	0.026
Cosine	93.65	0.016	79.27	0.033
Bray-Curtis	98.98	0.2	71.37	0.034
Squared Euclidean	97.94	0.016	75.29	0.026
Canberra	98.13	0.2	72.86	0.028
Manhattan	97.97	0.27	70.06	0.058

Fig.5.10 Table displaying the accuracy and time results for performing k-NN and k-NC on the training and validation data.

Overall, it is clear from the table in Fig. 5.9 that k-NN outperforms the k-NC approach, and both perform much better than the supervised SOM, with the highest accuracy of 98.98% provided by using k-NN with the Bray-Curtis similarity. The Bray-Curtis similarity can be calculated by taking the absolute difference of two vectors and dividing by their absolute summation, where for vectors, x and y, it can be expressed as:

$$BC = \frac{\sum(|x - y|)}{\sum(|x|) + \sum(|y|)}$$

A percentage accuracy this much higher than the supervised SOM clearly shows a vast improvement in the ability of k-NN to perform classification on reflectance spectral imaging data. However, there may be other alternative methods which can provide the same or better accuracy, and so these should be tested as well.

5.3.4 Support Vector Machine

The next classification model to be tested was the support vector machine (SVM). SVMs have demonstrated their use in performing classification in a wide variety of techniques and have seen use in spectral imaging applications for many years, including in remote sensing and cultural heritage studies (Gao et al., 2023; Polak et al., 2017) . SVMs can come in a variety of different forms and can use different parameters and inputs to generate the 'hyperplanes' which are used to cut data into separate classes in feature space. However, the main two parameters often adapted for different applications are the C value and the kernel. The C value is used to define how "smooth" a boundary between classes in feature space can be, where training data that is positioned on the wrong side of a hyperplane boundary can still be associated with the correct class, however the datapoints will incur a "cost", C, which can be exploited as a threshold to determine correct classification. The kernel instead seeks to project feature space into higher dimensional space before applying a hyperplane (known as "the kernel trick") allowing for more complex separation between classes which has shown to improve accuracy in many applications. To understand both how well SVM would perform in general on VNIR spectral imaging data recorded by PRISMS, but also to determine the optimal parameters for performing classification, an investigation was carried out to calculate the accuracy of the model while varying these different parameters. For this investigation, a C-Support Vector machine implemented in python using the scikit-learn module was tested using linear, polynomial, and radial basis function (RBF) kernels and different C values, with the accuracy and training time being monitored throughout the different parameter variations. A table showing the performance of the SVM using the different kernels can be seen in Fig. 5.11.

Support Vector Machine			
C Value	Kernel	SVM-C	
		Accuracy	Train Time
0.01	RBF	95.21	194.67
	Linear	92.06	204.45
	Poly d=2	94.32	148.31
	Poly d=3	94.86	58.57
	Poly d=4	94.9	56.62
0.1	RBF	97.54	59.62
	Linear	95.17	68.11
	Poly d=2	95.76	37.33
	Poly d=3	96	37.63
	Poly d=4	95.89	85.66
1	RBF	98.14	21.91
	Linear	95.77	29.19
	Poly d=2	96.24	32.61
	Poly d=3	96.35	78.31
	Poly d=4	96.26	480.08
10	RBF	98.34	12.83
	Linear	95.99	22.29
	Poly d=2	96.52	65.11
	Poly d=3	96.48	437.88
	Poly d=4	96.45	3160.9

Fig.5.11 Table showing the accuracy results of using a C-Support vector machine under various parameters. The best performing setup can be seen highlighted in green, which uses a C value of 10 and the RBF kernel.

As highlighted, the RBF kernel performed best out of the different approaches, however it was unknown if its accuracy could be increased even further. In response, the SVM using the RBF kernel was tested for different C values from 0.01 to 10,000, whereby it was found that the SVM was underfitted with C-values lower than approximately 10, and overfitted with values at and over approximately 1000 (Fig. 5.12).

Kernel	C Value	Accuracy	Train Time	
RBF	0.01	95.21	194.67	Underfitting
	0.1	97.54	59.62	
	1	98.14	21.91	
	10	98.34	12.83	Optimal
	100	98.44	13.09	
	1000	98.42	21.88	
	10000	98.39	69.83	Overfitting

Fig.5.12 Table showing the performance of the C-SVM using different C values with the RBF kernel. Underfitting setups can be seen in yellow, overfitting in red, and the optimal parameters can be found highlighted in green.

To estimate the true optimally fitted C value, different values ranging from 10 to 10000 were used in the training of the SVM when using the RBF kernel. In this test, the most optimal parameters will exist when the highest accuracy is reached, a graph showing the results of this investigation can be seen in Fig. 5.13 and shows a C value of 300 to be the most accurate at 98.444%.

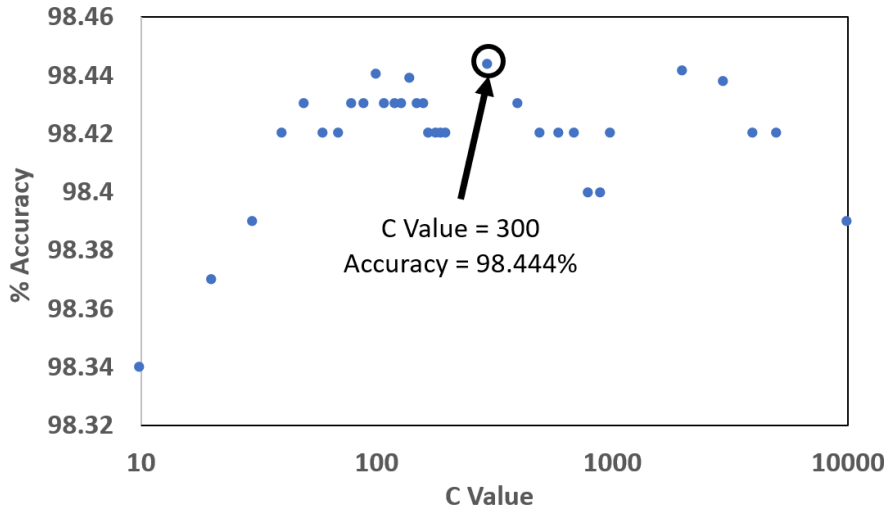

Fig.5.13 Graph showing the varied performance of the C-SVM under different C-Values ranging from 10-10,000. A C value of 300 was found to be the most accurate with a value of 98.444% and is shown in the graph by the arrow and circle.

5.3.5 Multilayer Perceptron (MLP)

Artificial neural network (ANN)-based classification approaches have seen more common use in recent years due to the rise in computational efficiency and ever-increasing amounts of big data. Some of the common ANNs can include convolutional neural networks, self-organising maps, or feedforward neural networks, each with varying performances for the classification of different types of data. One of the most implemented feedforward neural networks is the multilayer perceptron (MLP), which has proven to be able to classify spectral imaging data in many different studies, making it a good candidate for classifying the 14 classes of spectra extracted from the Javanese manuscripts. To evaluate its performance on VNIR PRISMS data, the MLP was trained and tested on the same data used in the previous tests, with the accuracy and training time being monitored for each variation of the technique. Optimisation parameters were tested by changing the activation functions used and investigating how the accuracy changed under the different settings, a table showing the results can be seen in Fig. 5.14. As highlighted in the table, the best MLP parameter to use for the reflectance data was the ‘relu’ activation function (see Eqn. (2) of (Fukushima, 1975) with an accuracy of 98.28%. While not as accurate as the k-NN or SVM, an accuracy of over 98% on a set of validation data is good.

Multilayer Perceptron (MLP)		
Activation Function	Accuracy (%)	Train Time (s)
relu	98.28	77.87
identity	95.5	36.64
logistic	98.02	216.9
tanh	98.17	118.69

Fig.5.14 Table showing the performance of the Multilayer Perceptron using different activation functions. As hhighlighted in the table, the Relu activation function was found to provide the highest accuracy.

5.3.6 Decision Trees (DT) & Random Forest (RF)

The final techniques covered in this investigation were the random forest (RF) and decision tree (DT) models. The two models were trained and tested in the same way as all previous techniques, with the separation of training data and validation data being used to evaluate the models. However, for any decision tree classifier different criteria can be used to define the quality of a split in the tree, and so variations of these were also investigated to examine if the techniques can be optimised for spectral imaging data, thereby providing higher a classification accuracy. The results from investigating both DT and RF can be seen for three different criteria, Gini, Entropy, and Log Loss, can be seen within the table in Fig. 5.15, with the best performing of each model highlighted in green.

Criterion	Decision Tree		Random Forest	
	Accuracy (%)	Training Time (s)	Accuracy (%)	Training Time (s)
Gini	96.57	4.06	98.22	82.67
Entropy	96.82	4.22	98.3	98.82
Log Loss	96.83	4.29	98.28	98.94

Fig.5.15 Table showing the performance of both the decision tree and random forest classification models using different classification criteria. The best performing methods can be seen highlighted in green.

The results from Fig. 5.15 show that both the decision tree and random forest perform well in performing classification on the spectral reflectance data, with random forest slightly outperforming the decision tree approach by 1.99% at the cost of taking significantly longer to train, though a training time of 98.82 seconds is small regardless. With an accuracy of 98.3% it would be interesting to further examine the random forest model alongside the other best performing techniques to investigate which will continue to perform, and which will perform best, for datasets larger than the sample data of 20,000 spectra already used in training and validation.

5.3.7 Cross Evaluation of the Best Performing Models

When applying classification models on large collections, the size of the datasets being analysed will be significantly larger than those used in the previous evaluations of the different techniques. Therefore, it would be wise to confirm which is the best performing classification technique when using much greater sample sizes, so that a better representation of the final target data, which is more likely to be encountered in practical terms, can be used. In response, the same classes were used for training the different models as before, but with the training and test data changing from a 70:30 split of 20,000 samples per class to a 10:90 split for 200,000 samples per class respectively, meaning that each model would now be trained using 20,000 samples per class and validated using 180,000 samples per class.

The four techniques that were confirmed as best to use during the previous tests were:

- K-NN with 9 neighbours using the Bray-Curtis distance metric.
- C-SVM using a C value of 300 with the RBF kernel.
- Random Forest using entropy.
- Multilayer Perceptron with a Relu activation function.

Therefore, all these techniques were investigated for the new data, with their accuracies, training time, validation time and confusion matrices being recorded simultaneously. A table showing the results for this investigation can be seen in Fig. 5.16.

Model	Accuracy (%)	Time (secs)	
		Training	Validation
k-Nearest Neighbours	98.035	0.333	1327.294
Support Vector Machine	98.456	29.896	2008.604
Random Forest	98.315	145.941	31.208
Multilayer Perceptron	98.285	117.355	4.305

Fig.5.16 Table showing the large-scale validation results for the four best classification techniques. Highlighted in green is the C-SVM, which showed the best performance.

As highlighted in the table, most of the classification models performed similarly well, with the C-Support Vector Machine based classification method showing the highest accuracy percentage of 98.456%, but only higher than all the other techniques by at least 0.1%. This suggests the C-SVM is likely the best approach to use for classifying this specific large collection of spectra, however as the other techniques are not far behind in their accuracy, it is not guaranteed that the C-SVM would perform best as a general approach. It is interesting though that the previously best performing approach, k-NN, now performs the worst. It is important to note as well that the multilayer perceptron model was able to achieve a similar percentage accuracy to the C-SVM but

in a much shorter amount of total (training + validation) time. This could suggest that for other applications where classification is to be performed 'live', as in during data collection, this approach may be better to use.

Another interesting result to note about the SVM however is that the accuracy for the larger dataset is higher than seen for the previously tested smaller sample sizes, suggesting that even though the proportional size of the test data has decreased from 70:30 to 10:90 (training: validation), the overall increase in the number of training spectra has still resulted in an increase in accuracy. This agrees with a conclusion given in Foody & Mathur, (2004) which found that SVMs are positively correlated with training dataset sizes, implying that not only would SVM be useful in this specific scenario of analysing data using clustering results derived from large collections, but it may continue to perform even better as the amount of training data increases over time. In future, it would therefore be wise to further test this to investigate also how increases in VNIR spectral complexity may also affect the performance as training data also becomes larger.

While the SVM possessed the highest accuracy, each approach was still quite similar in its performance. To finally provide a more conclusive result for which classification model would be the most reliable, the confusion matrices were also examined to understand where misclassifications were distributed. The confusion matrices were plotted with their diagonals set to zero, so that the distribution of misclassifications could be visualised more clearly as a map. The corresponding maps are shown in Fig. 5.17.

Each of the confusion matrices appears to have the same structure, likely suggesting that the same misclassifications exist across all methods. However, it could also suggest that some of the validation data was labelled incorrectly before training or validating any of the classification models, and that each technique performs equally well given the data used for training. Either way, these maps confirm that there is no considerable difference between the techniques in terms of the distribution of incorrectly classified data, and therefore the C-Support Vector Machine using a C Value of 300 with the RBF kernel should be used for classifying the 14 pigment mixtures over the rest of the collection, as there is no significant structural change in the confusion matrix from any of the other techniques and it possesses the highest overall accuracy.

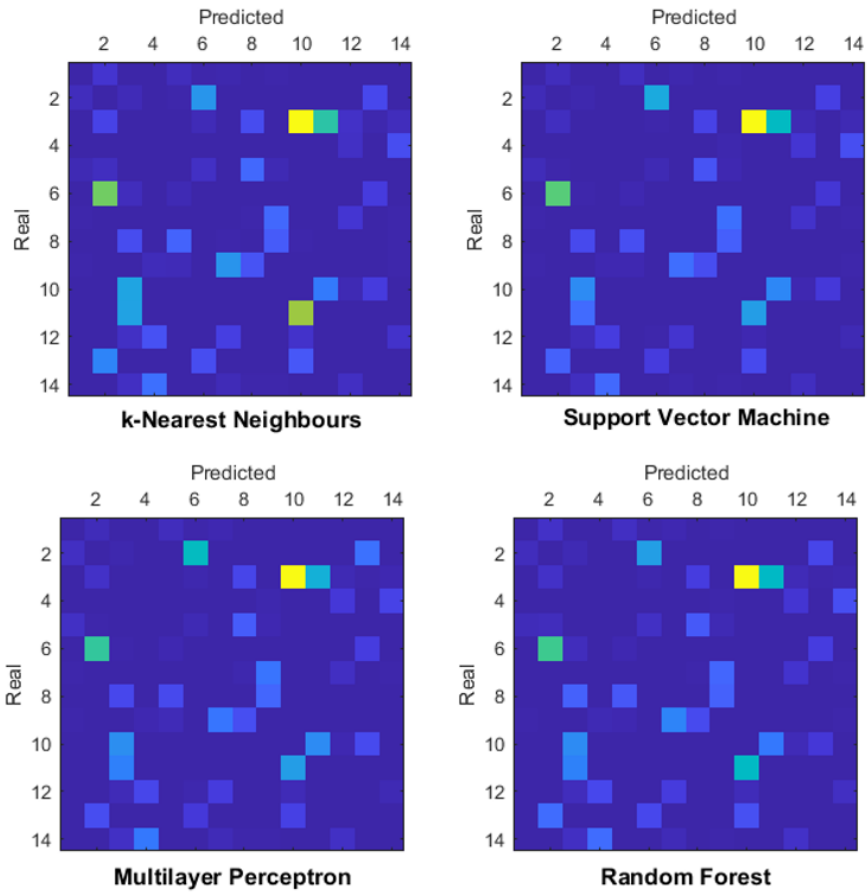

Fig.5.17 Confusion matrices for the four best performing classification techniques which show very similar structure.

5.4 C-SVM on the British Library's Maritime Southeast Asian Manuscript Collection

Now that it was established that the C-support vector machine was the most accurate method to be used when training and validating the different classification models for larger scale data, an investigation into how well the trained C-SVM could classify a set of target data covering the rest of the collection was performed.

Class & Probability Maps

To maintain continuity with the original and group cluster maps, new 'class maps' were built to represent the distribution of pixels belonging to a class where each pixel can be characterised by its given integer intensity which is equal to their class number. For example, a group of pixels belonging to class -3, corresponding to the mixture of vermillion and bone ash white, will all have an intensity of -3. This way, class maps can easily be compared between the original training/validation data and the newly introduced target data to determine how well the classification has performed.

An additional way of monitoring the performance of the classification procedure is through the creation of a probability map, where each pixel possesses a probability calculated through cross validation, which represents how well each pixel or cluster belongs to its designated class after being assigned by the C-SVM. Probability map pixels will always possess an intensity ranging from 0-1, with 1 showing a high probability that the pixel belongs to its assigned class. These probability maps were produced as they can be useful in checking how accurate pixel classification may be, but in addition to this they can also be used to remove any low probability pixels, thereby creating a 'high probability class map', where pixels with low values for probability will be given an intensity of 0, to indicate that they are unclassified. As SVM has no procedure for dealing with data that does not fit well into any of the trained classes, the production of these high probability class maps was necessary to ensure that many pixels with spectra representing unknown mixtures would not be incorrectly grouped into trained classes solely because they are forced by the algorithm to belong to a class.

High probability class maps are produced by first determining what threshold probability value should be used to decide when to remove pixels or clusters from their assigned class. This is done by computing the mean probability and standard deviation for all the correctly assigned validation data per class, which can then be used to calculate cut-off points equivalent to either the mean, or the mean subtracted by a certain number of standard deviations. Calculations for the mean and standard deviations of the probabilities can be seen within the table in Fig. 5.18.

Class	Probability		Population Percentage Above Mean - N Std Dev (%)			
			N Standard Deviations			
	Mean	Std Dev.	0 std dev	1 std dev	2 std dev	3 std dev
All	0.9874	0.0495	86.355198	95.20113208	96.913836	97.78812352
-14	0.997622	0.024205	96.1760201	98.61077844	99.0179641	99.2240635
-13	0.97732	0.063114	81.8704955	92.87583731	95.7790905	97.23354343
-12	0.982032	0.066208	88.2507662	94.09309207	95.9507063	97.09850522
-11	0.991182	0.039289	88.518997	95.85661471	97.3813	98.12294802
-10	0.994851	0.032779	93.0830946	97.37433981	98.2230776	98.6247874
-9	0.988206	0.052524	90.2263848	95.50016776	96.9206885	97.67416733
-8	0.991301	0.037606	87.7996388	96.12951351	97.521565	98.21927335
-7	0.989824	0.045957	89.5443402	95.66714572	97.1319075	97.88091414
-6	0.987392	0.047494	86.993613	95.23423652	96.9927689	97.82729805
-5	0.961132	0.070096	79.2740905	91.64863253	95.2111556	96.96288092
-4	0.980308	0.066271	86.0522249	93.73765277	95.9486921	97.16572606
-3	0.993973	0.035685	91.9867116	97.01906568	98.0039485	98.46087594
-2	0.994052	0.037521	93.8284872	97.11032093	97.9300761	98.42688032
-1	0.993882	0.034112	90.7112928	96.96490372	97.9941777	98.491146

*Fig.5.18 Table showing the mean, standard deviation, and population of correctly labelled data above the mean-N*Standard deviations for all 14 classes.*

As the table shows, most correctly assigned data will belong within 3 standard deviations of the mean for each class, with a small amount of leeway to remove probably incorrectly classified spectra. It was therefore decided that the probability threshold used would be the mean subtracted by 3 standard deviations, where, for example, the threshold, P, for class -14 would be:

$$P = Mean_{class14} - 3 \times Standard\ Deviation_{class14}$$

$$= 0.997622 - 3 \times 0.024205$$

$$= 0.925117$$

Therefore, for class -14, any pixel with a class probability of less than 0.925117 would be reclassed as unassigned and be given an intensity of zero in the high probability class maps. This procedure was chosen because the probability for a pixel to belong to its class varies depending on the class and its variation of constituent spectra, as can be seen in Fig. 5.18 where the mean probability can fluctuate significantly between classes. Therefore, by considering the probability with which correctly assigned pixels belong to their class during validation, a reasonable judgement can be made to determine if any new data belongs as well. This results in much 'cleaner' class maps, where only the spectra which are similar in nature to the 14 classes are assigned, and other pixels or clusters with completely different and untrained spectra will be ignored, allowing for more accurate classification and more clear inspection and analysis. It is important to note however that these unclassified spectra are not completely removed, and are typically given a class value of zero, so that they can be inspected or perhaps clustered by themselves at a later stage. This means that in the grand scheme of analysing a new collection using original data from an already characterised previous collection, data analysis can be performed in a loop to continually update databases and perform pigment mixture identification over multiple collections.

5.4.1 Pixel Level Classification using the C-Support Vector Machine

The approach investigated involved performing classification on each individual pixel in every spectral imaging data cube recorded using PRISMS in the VNIR, resulting in the creation of a set of 1515 high probability class maps which illustrate the distribution the 14 different pigment mixtures trained into the SVM. This pixel classification method involved taking the trained SVM, using the RBF kernel with a C value of 300, and using it to classify the spectral reflectance data recorded for the manuscript collection. To reduce the computational time for performing the classification, each spectral imaging data cube was resized by a factor of 0.25 to possess $1/16^{th}$ of its original number of reflectance spectra. This massively reduced the computational and time

requirements for performing classification for all data cubes, while still allowing for most small spatial regions with different spectral information to be grouped into their respective classes correctly. Following this, each of the resized data cubes were used as an input into the trained SVM, producing a class map and probability map which were used together to produce the high probability class maps. As mentioned in the previous section, high probability maps were produced for the different standard deviation ranges about the mean, so that pixels with a low probability of belonging to their class become unclassified. The process followed for pixel level classification is demonstrated in the flow chart shown in Fig. 5.19.

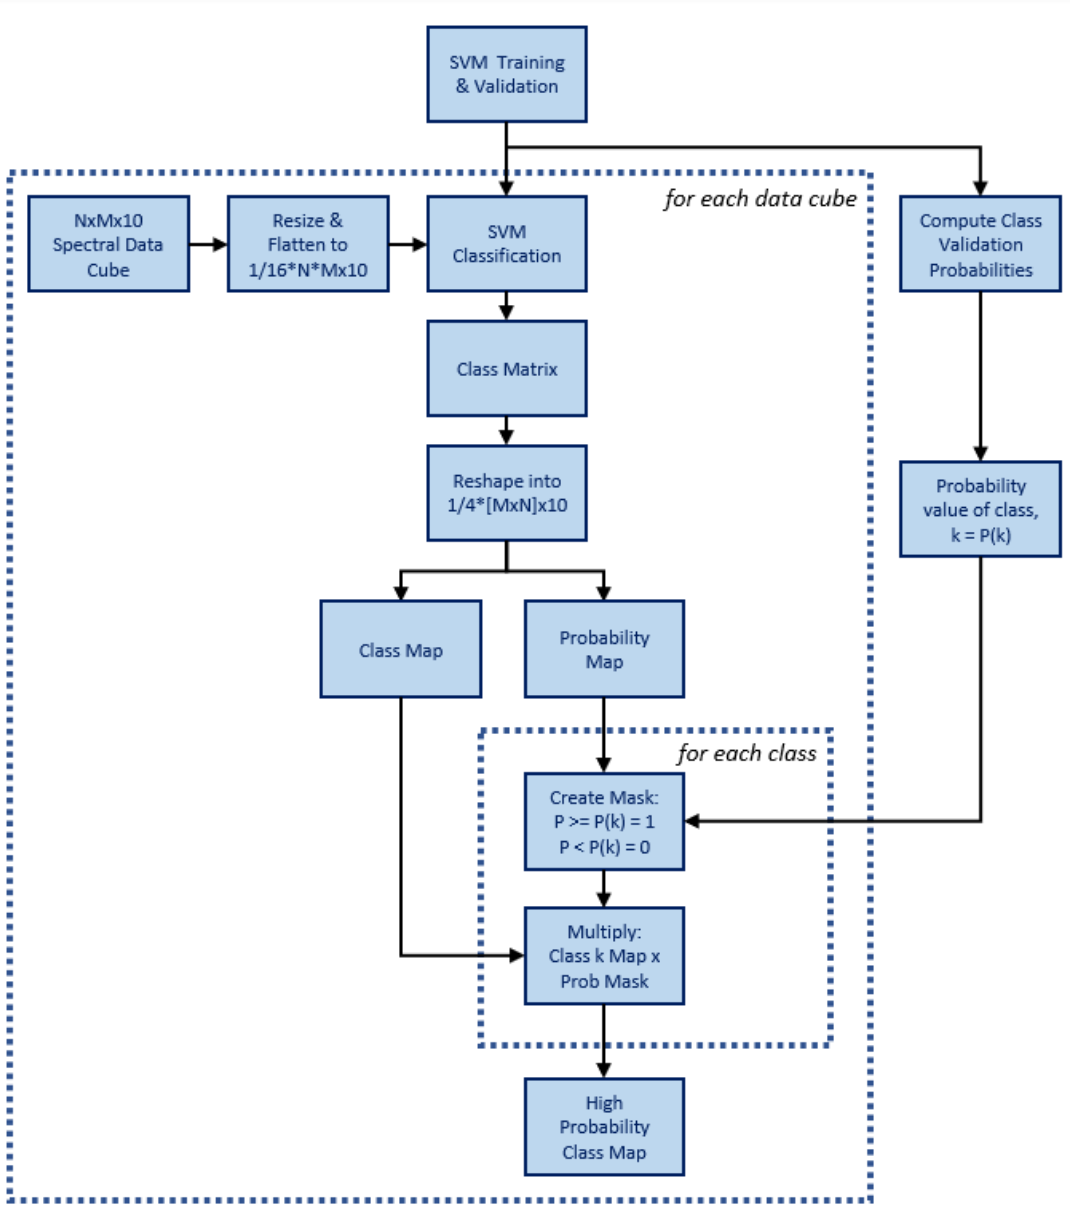

Fig.5.19 Flow chart showing the process of performing C-SVM classification to produce high probability class maps for all spectral data cubes within a collection.

5.4.2 Pixel Level Classification Results & Discussion

Computational & Time Requirements

Probably the largest hurdle with approaching classification in an image-by-image way for large collections are the computational and time requirements, which are quite high considering that millions or even billions of spectra need to be classified by the SVM. When performing the classification on a computer with an AMD Ryzen 5 5600G CPU (12 CPUs ~3.9GHz), with 16384MB RAM, each of the resized spectral imaging data cubes took a mean average of 117.6 seconds to be opened, classified and stored as class and probability maps. The total runtime of performing the classification for all 1515 reflectance spectral data cubes was ≈49.5 hours, which is considered a large amount of time when taking into account that this data had already been reduced in size by a factor of 16. Assuming the performance of the SVM and computer remains stable and is linear over time, this would imply that it may take approximately 33 days to complete the pixel level classification of a collection of 1515 reflectance spectral data cubes without any resizing.

It is therefore important to emphasise that this amount of processing time is much larger than that of performing clustering and suggests that performing classification on the clustering results in the future may therefore be more beneficial in terms of time efficiency, so long as the accuracy of the clustering results themselves is reliably high to avoid misclassification. What has been shown in the previous chapters of this thesis however is that this is normally achievable by using fine parameters during automated clustering, and as the extra clusters likely to be created due to the increased precision would be dealt with using classification models instead of grouping, the workload required for their analysis would be greatly reduced. This method should therefore be explored in the future, however in the meantime it was also important to understand if the classification model could perform in a satisfactory way by analysing cluster maps directly, as a cluster is an accurate representation of real spectra anyway.

Classification Performance for all Manuscripts

To evaluate the pixel level classification performance of the C-SVM, each manuscript was inspected to determine whether the pixel assignments led to each class falling into one of 7 evaluation groups, which were as follows:

- ‘Mostly True Positives’ – meaning that pixels possessing spectra corresponding to the correct pigment mixtures were the majority in the class. Although some leniency was permitted for small groups or sparse pixels with incorrect spectra, as it was not uncommon for random incorrect pixels to still have a high enough probability to not be unclassified.
- ‘Mostly True Positives with False Positives’ – meaning that the manuscript mainly follows the previous ‘Mostly True Positives’ evaluation group, but there are also other

pigment mixtures which have been assigned to the class, normally due to them possessing similar spectra.

- ‘Mostly False Positives with True Positives’ – where most pixels assigned to the class are not of the correct pigment mixture but may be assigned due to their similarity in spectral features. For example, where chrome yellow may be placed into the orpiment class. In some of these manuscripts there also may be small areas of true positive classification, but they are the minority.
- ‘Mostly False Positives’ – where most pixels assigned to the class are not the correct pigment mixture but are normally assigned due to their similar spectral features. This normally occurs when the pigments differ from those seen in the four Javanese manuscripts used for training.
- ‘Mostly False Negatives with True Positives’ – meaning that there are very few pixels in areas where the correct pigment mixture exists, but there are some true positive assignments. This illustrates that the classification works for many pixels but perhaps requires more training data to cover a wider variety of potential spectra due to changes in pigment concentration.
- ‘Mostly False Negatives’ – where there are no assigned pixels in areas where the correct pigment mixture exists, this illustrates a total failure of the SVM to perform classification but perhaps could be addressed by using more training data which can account for a wider variety of spectra.
- ‘Pigment Not Present so N/A’ – meaning that no pixels were ever assigned to the class in the manuscript, aside from a few randomly associated spectra. In these scenarios, this is not a problem however, as there is no evidence for the pigment mixture to exist on the manuscript, essentially making this an evaluation group for ‘True Negatives’.

After analysing every manuscript for which PRISMS data had been recorded, the different evaluation groups were assigned to each class for every manuscript, the results of which can be seen in Fig. 5.20. What is clear from this table, is that most classes throughout all the different manuscripts fall into either ‘Mostly True Positives’, ‘Mostly True Positives with False Positives’, ‘Mostly False Positives with True Positives’, or ‘Pigment not Present so N/A’. These results suggest that classification using SVM performs well on new data but requires a greater number of classes with different mixtures to use during training if false positives are to be avoided in the future. The reasoning behind this is that most false positives occur as a result of pixels with an unseen and untrained mixture existing on a manuscript which happens to have similar spectra to the class with which it is assigned. Some of examples of these false positives, along with some true positive and false negative data are discussed in the following parts of this section. In these sections some of the limitations of the approach taken to perform classification are demonstrated, with some interesting results covered which open new questions for future research. The sections

are mainly covered by class but start with the areas that possess false negatives, as these show the performance of the SVM at its worst.

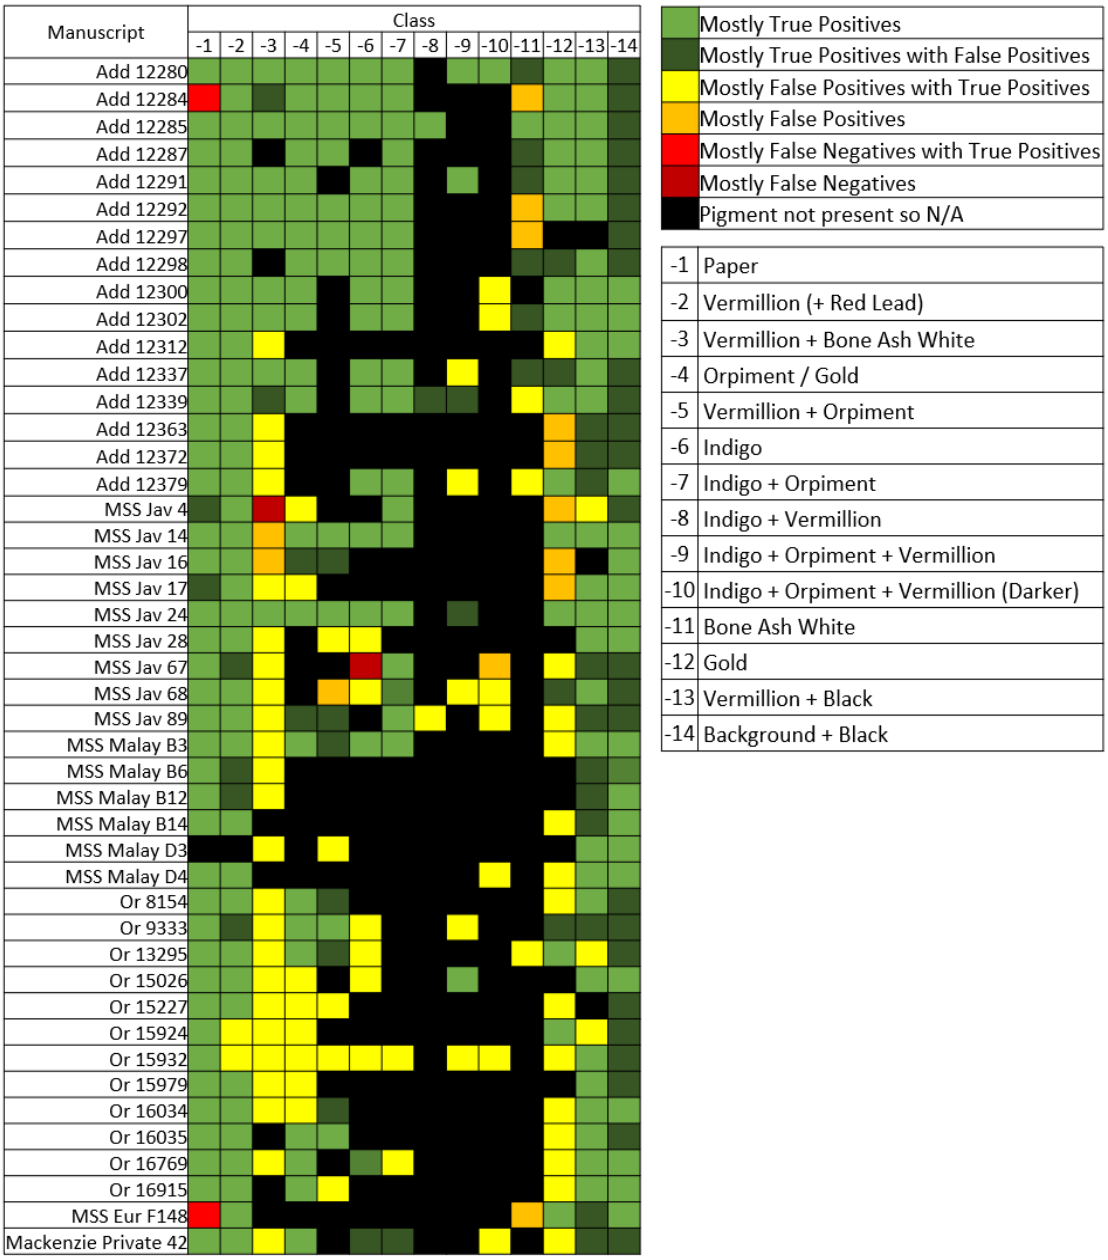

Fig.5.20 Tables showing - Left: the classification evaluation results for the 14 classes over the different manuscripts within the collection. Top Right: The colour key for the results. Bottom Right: Which material or pigment mixture the class represents.

5.4.2.1 Class -1 (Paper)

The classification for the paper is largely performed well, with most pixel assignments being true positives with some false positives included in two manuscripts. However, there are two scenarios where the paper is completely misclassified, and both times this is done by mistakenly grouping the paper into class -11, for which real colour class maps can be seen in Fig. 5.21.

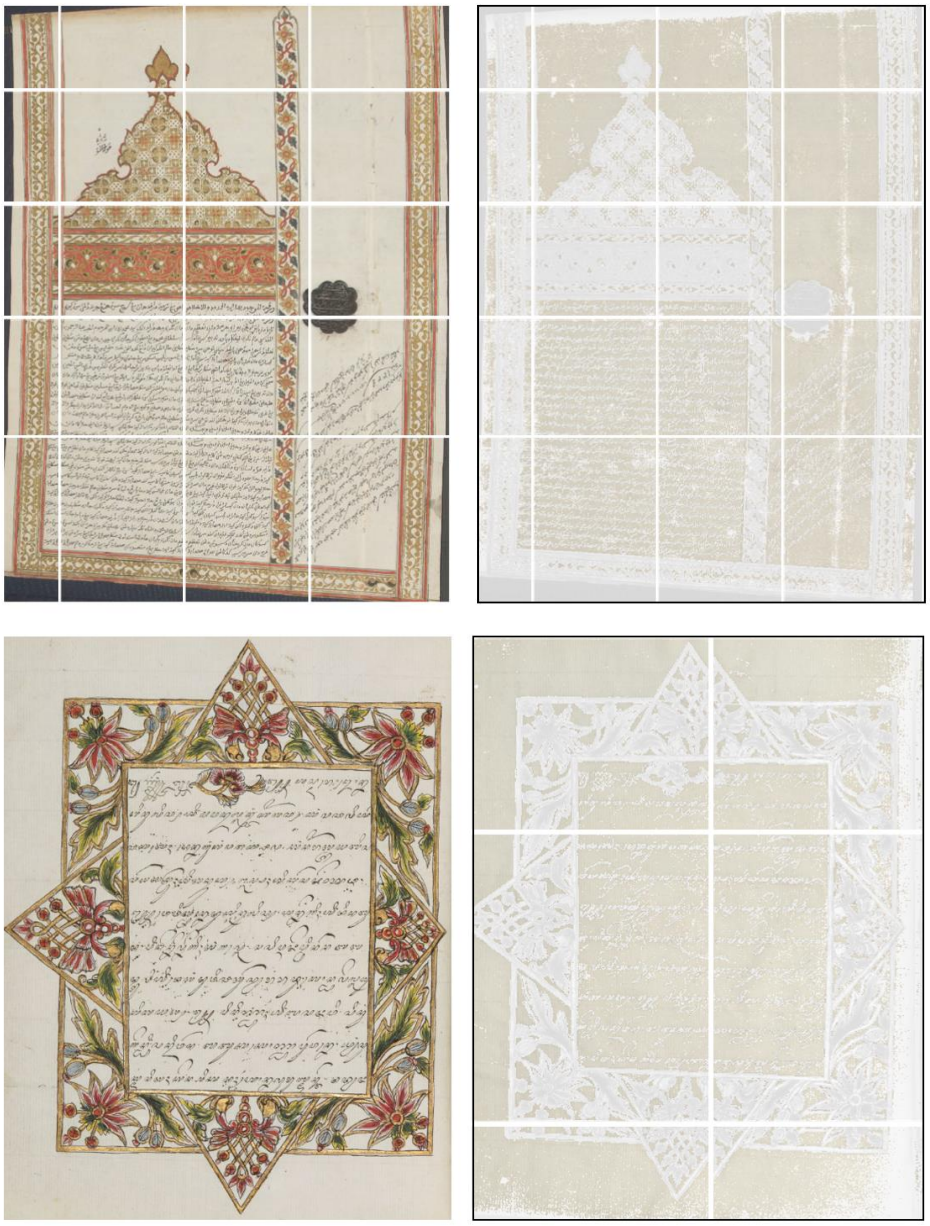

Fig.5.21 True colour high probability class maps representing the incorrectly classified pixels belonging to class -11 (White) for manuscripts Add MS 12284 and MSS Eur F/148/4 f106.

The two regions where the C-SVM fails are in Add MS 12284 and MSS Eur F/148/4 f106, where both times the paper is assigned to the bone ash white class due to the spectra having a lower intensity towards the near infrared than that seen for most other paper spectra in the collection, as

shown in Fig. 5.22, where the paper spectrum from Add MS 12280 f3r, which was used to train the SVM, can be seen plotted alongside the misclassified paper spectra.

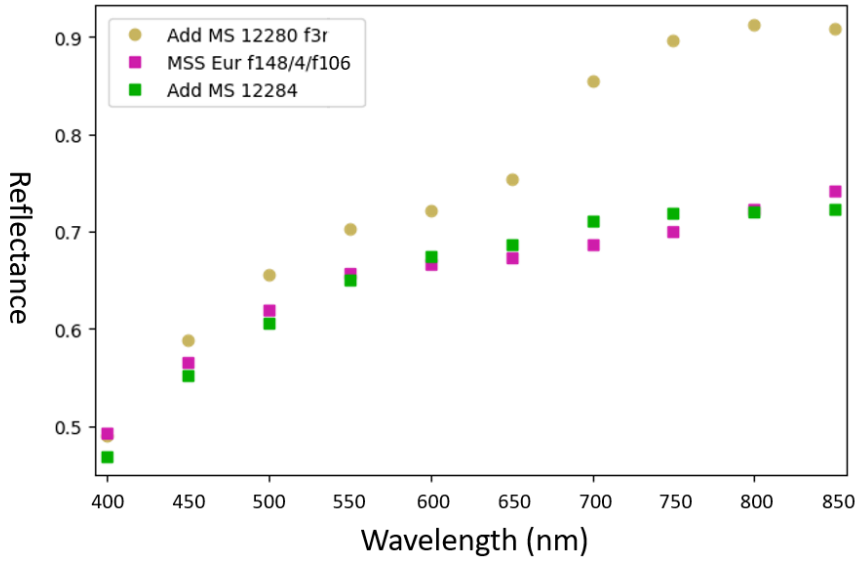

Fig.5.22 Graph showing the PRISMS reflectance spectra extracted for the paper from three different manuscripts. Add MS 112280 f3r clearly demonstrates a spectral difference towards the NIR (>650nm).

The similarly low intensity in both manuscripts could be related to them both being produced using English paper, with MSS Eur f148/4 f106 possessing a watermark 'G. Jones 1807' and Add MS 12284 watermarked with 'S & Wise 1815'. Furthermore, on close inspection of the paper in both manuscripts, straight lines indicative of marks commonly found on laid paper can be seen, a feature not seen on many other manuscripts, suggesting that they may also share a similar manufacturing process which has resulted in this difference in reflectance.

5.4.2.2 Class -2 (Vermillion Mixtures)

Throughout the entire collection the main red pigment to be found was vermillion, and therefore it would be wise to assume that the class representing vermillion would be found in many other manuscripts. After performing classification, large numbers of correctly classified vermillion pixels could be found over many different manuscripts, most of which were confirmed by cross referencing results seen in chapter 3, where more in depth analysis of the British Library's maritime southeast Asian collection was performed. The confirmation of successfully classified red vermillion mixture regions can be seen in the class chart shown previously (see Fig. 5.20) however a few example manuscripts and pages from different regions in maritime Southeast Asia are also covered in closer detail in the following subsections to demonstrate how the SVM is able to perform classification in manuscripts from regions separate from Java (the location of the 4

training manuscripts). The manuscripts/pages explored further include Or 8154 f3v, Or 15227 (Multiple Pages), and Or 16034 f261r, and were chosen as they share provenances with the different regions which possessed the greatest abundance of manuscripts outside Java including Sumatra, Sulawesi and also the Malay Peninsula.

Or 8154 f3v

In the Buginese manuscript Or 8154 f3v, the vermillion class was detected in the red wings and other red decorative parts of the illustration of Buraq, the winged horse (see Fig. 5.23). In chapter 3 these regions were confirmed to be vermillion using both XRF, which showed presence of mercury (Hg) and sulphur (S), and Kubelka-Munk mixture modelling which showed that the mixture of paper and vermillion fitted to the red spectra well.

Fig.5.23 RGB Image (Left) and true colour class maps (Right) for the correctly classified vermillion pixels seen in manuscript Or 8154, on folio f3v.

These results are particularly interesting as it implies that even if the paper substrate used below vermillion is different, or if a different artistic style is used in a different location, classification can still be performed accurately with the SVM.

Or 15227

The vermillion class can also be found in several pages of manuscript Or 15227, a copy of a Qur'an from the Malay Peninsula. In this manuscript, pixels were assigned to the class both in the decorative frames seen throughout various pages and in the many pages where juz' are marked with illuminated roundels and floral patterns, examples of which can be seen in Fig. 5.24.

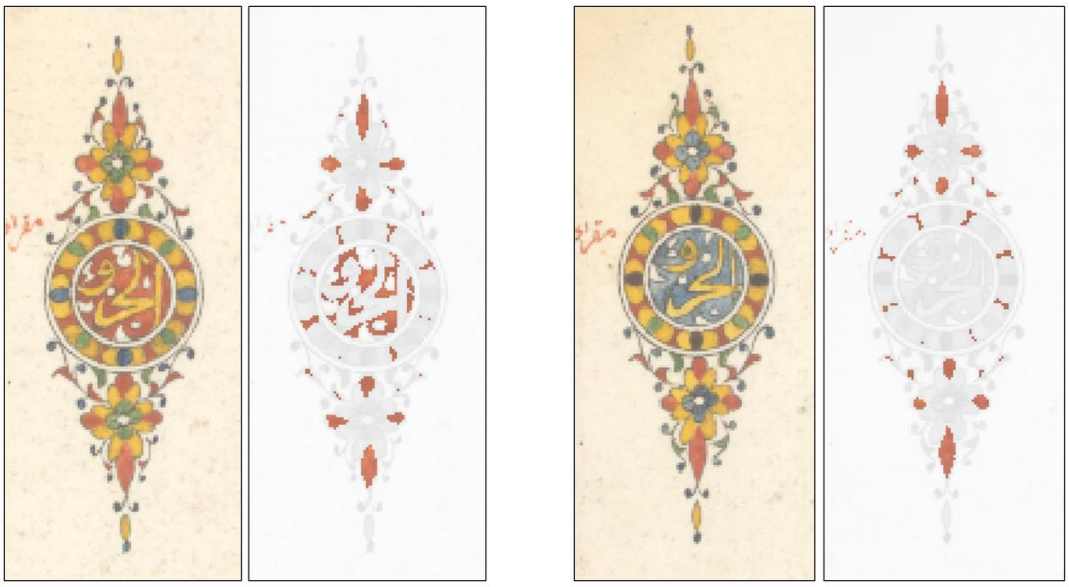

Fig.5.24 True colour class maps representing vermillion classified pixels for juz' markers within Or 15227

Or 16034 f261r

A further manuscript from Aceh can be seen with classified vermillion pixels also, where, like many others, they can be found distributed in patterns over the opening double frames, as demonstrated in Fig. 5.25.

Fig.5.25 True colour class maps showing classified vermillion pixels in manuscript Or 16034, folio f261r.

5.4.2.3 Class -3 (Vermillion + Bone Ash White Mixture)

For the vermillion and bone ash white class, correctly classified pixels can be seen in many of the Javanese manuscripts where the mixture was confirmed in chapter 3. A good example to illustrate this is in Add MS 12292 f2v, where the pink can be seen in the opening double frames positioned in both the rubrication aside the text and in the flower petal motifs placed around the outsides (Fig. 5.26).

Fig.5.26 RGB Image (Left) and true colour class maps showing the distribution of pink vermillion + bone ash white pixels in manuscript Ass MS 12292, folio f2v.

In all the other Javanese manuscripts where the pink vermillion and bone ash white pigment mixture existed, it was found to be successfully assigned to class -3, again illustrating the usefulness of SVM in performing classification. There were, however, also many circumstances for this class where pixels were incorrectly assigned due to them possessing similar spectra but not the same materials, for example, where red dye or a thinly applied vermillion pixels are classified as the same mixture. This is no surprise considering that most of the training data used for the vermillion areas in the training data set were applied heavily and would therefore not be trained to fit thinly applied regions, and likewise as no dyes were trained into the SVM either, misclassifications between the pink and a spectrally similar dye are likely to occur. A good example of the SVM incorrectly classifying a dye-based pink into the vermillion-based pink class can be seen in MSS Malay B12 f2r, a high probability class map of which can be seen in Fig. 5.27.

Fig.5.27 Images showing the RGB (Left), and true colour class map (Right) of false positive red dye pixels assigned into the pink vermillion and bone ash white class due to their similarity in VNIR spectra.

5.4.2.4 Cass -4 (Orpiment)

Throughout the collection, numerous different regions with spectra corresponding to orpiment were correctly classified into class -4 using the SVM. Orpiment can be found correctly classified in many other parts of the target data within the collection and covers a variation of provenances including manuscripts from Sulawesi, the Malay Peninsula, Aceh, and other parts of Java, with some good examples including Or 8154 and Or 13295.

Or 8154 f3v

In the Buginese manuscript Or 8154 in folio f3v, orpiment can be found correctly classified where the yellow body of the winged horse is placed in the orpiment class (see Fig. 5.28). While the majority is properly classified, the front legs of the horse are not placed into the high probability class maps however, and this appears to be due to a darker more degraded look to the paints.

Or 13295 f190v

In the case of manuscript Or 13295 f190v, many of the yellow pixels on the outer part of the frame have been successfully classified into the orpiment class as well, as demonstrated in Fig. 5.29.

Fig.5.28 Images showing the RGB image for Or 8154 f3v (Left), with the true colour class maps for orpiment (Right) created from the C-SVM.

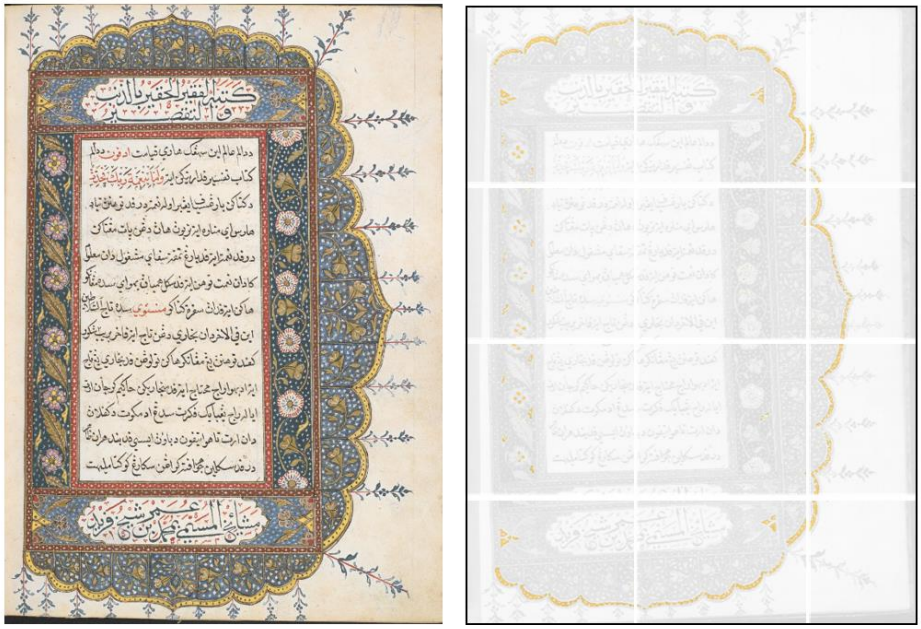

Fig.5.29 Images showing the RGB image for Or 13295 f190v (Left), with the true colour class maps representing orpiment pixels (Right)

False Positive Misclassifications

The presence of orpiment on many of these other manuscripts can be confirmed by point measurement analysis using techniques such as XRF or Raman spectroscopy, however there are some regions in other manuscripts which possess misclassifications where a yellow spectrum has been assigned into class -4 but is not orpiment. One such manuscript where this happens regularly can be seen in Or 15227, where a chrome yellow-based pigment which is present in many of the decorative frames and juz markers can be seen incorrectly classified with the orpiment (Fig. 5.30).

Fig.5.30 Images showing the RGB and true colour class maps for a juz' marker and an opening frame in manuscript Or 15227. In both images, pixels corresponding to chrome yellow can be seen classified into the orpiment class.

The misclassification of yellow pigments in Or 15227 comes as no surprise, as it can be incredibly difficult, and sometimes impossible, to distinguish between pigments such as orpiment and chrome yellow, through the use of VNIR spectral imaging alone. It is often therefore a requirement that XRF or other techniques such as Raman spectroscopy need to be used to fully distinguish between them, so without this information being factored into the classification in some way, such a misclassification will always be likely to occur.

This does however introduce a research question that would be very interesting to explore in the future, which is whether classification can be performed multimodally. This is because, by using multimodal clustering, as seen in chapter 4, to create sets of training data with information from other sources such as XRF spectral imaging, or perhaps SWIR reflectance, there is potential for more accurate holistic classification to be performed on new collections. Such an approach, for example, would be likely able to distinguish between yellow pigments of different kinds, where, for example, XRF would be able to provide the elemental composition of different painted regions which would either confirm or deny the presence of specific materials such as arsenic-based or chrome-based yellows.

5.4.2.5 Class -5 (Vermillion + Orpiment)

When investigating the classification of the orange orpiment and vermillion mixture class, most manuscripts primarily have true positive pixel assignments. One good example of this can be seen in the Javanese manuscript Add MS 12287, where the use of orange in the opening double frames is classified into the orpiment and vermillion mixture class correctly (Fig. 5.31).

While most of the data throughout the collection is correctly classified, there are some circumstances where there are large inclusions of false positives due to a lack of training data covering dark yellow pigments. For example, misclassifications can be seen on manuscript Or 15227, where some of the darker chrome yellow illumination in the decorative frames is assigned incorrectly due to its similarity in spectral shape and intensity to the training data used in the SVM for class -5. This can be seen in Fig. 5.32, where the dark yellow pixels in the outer frames are seen to be assigned into the same class as the orange found on Add MSS 12285 f1v, which was used in the training data set.

Fig.5.31 Images showing the RGB (Left) and true colour high probability class maps for the orange vermillion and orpiment mixture seen in Add MS 12287.

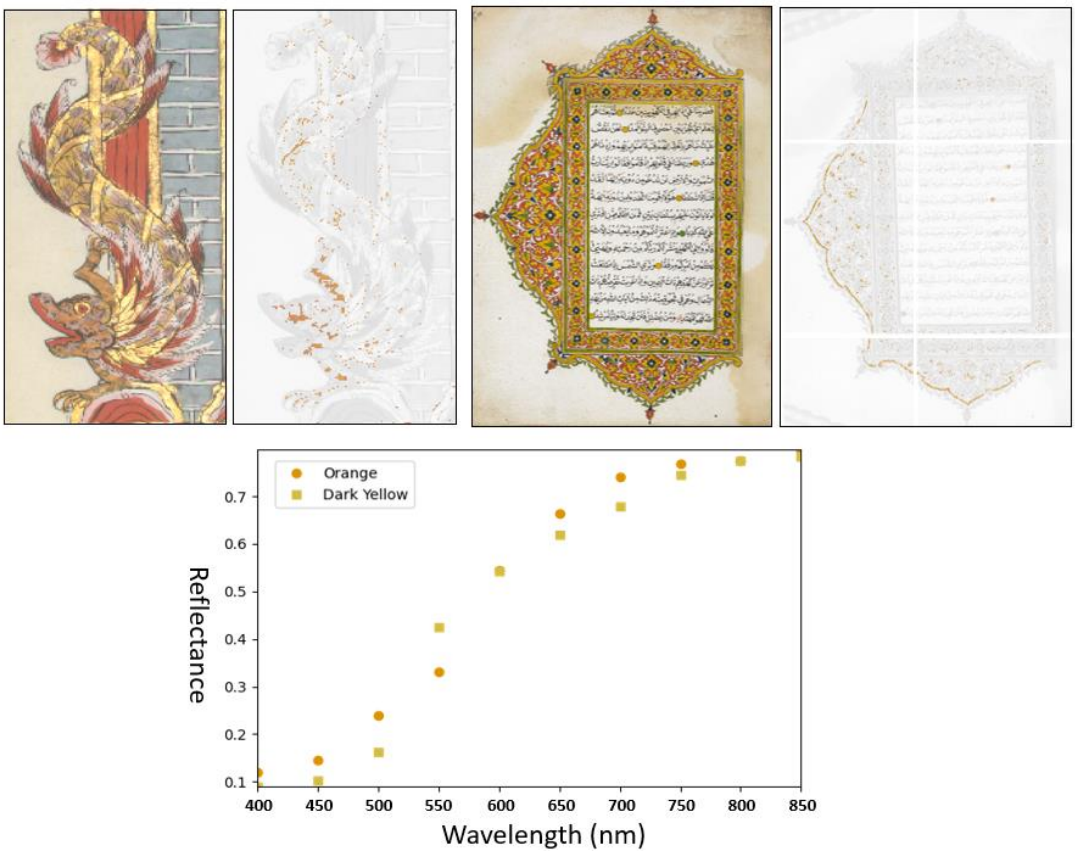

Fig.5.32 Upper Left - Images showing the RGB (Left) and true colour class map of an orange used for training (Right). Upper Right – Images showing the RGB (Left) and true colour class map of an incorrectly classified dark yellow in Or 15227. Bottom – extracted VNIR spectra from the two areas indicating their similarity in spectrum which is suspected to be the root cause for the misclassification.

5.4.2.6 Class -6 (Indigo)

In most of the manuscripts in which indigo existed in the collection, the C-SVM was able to correctly classify most of the pixels into the proper class, as demonstrated in the classification table (Fig. 5.20). However, there is one manuscript which completely failed to identify indigo, as is noted in the class table for MSS Jav 67. In MSS Jav 67 f12r, the blue used for colouring the illustration of a man's robes is indigo, but there are very few pixels which are properly classified into the indigo class, with the majority instead being assigned to the lilac and green indigo-based mixtures, as shown in Fig. 5.33. It is likely that the assignment into these alternative classes is mainly due to changes in the thickness of the indigo applied onto the paper causing changes in the reflectance spectrum, where the darker and more yellow looking paper introduces spectral features which fit more closely to a green or lilac spectrum. For example, if the spectrum of the green area taken from MSS Jav 24 during training is plotted next to the MSS Jav 67 indigo spectrum incorrectly assigned to class -7, then it becomes clear that the two spectra are very similar, and likewise the same problem is detected for class -8 (see Fig. 5.34).

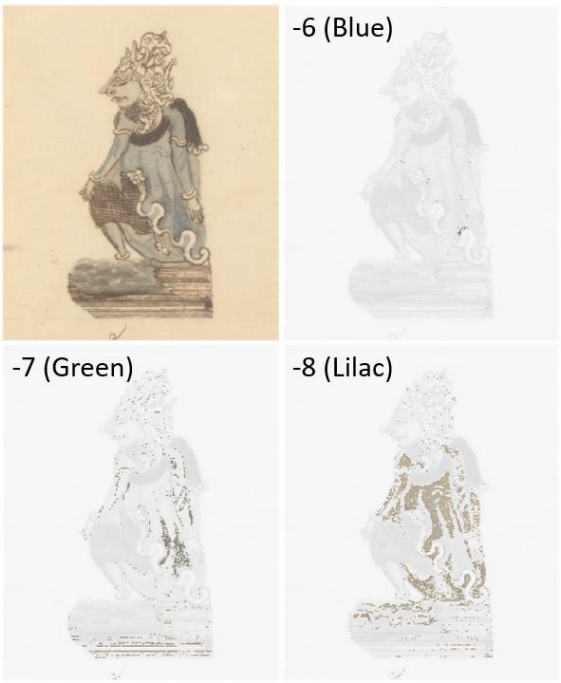

Fig.5.33 RGB (Top Left) and true colour class maps for classes -6, -7, and -8, which show that most pixels are incorrectly assigned to the lilac and green classes on MSS Jav 67 f12r.

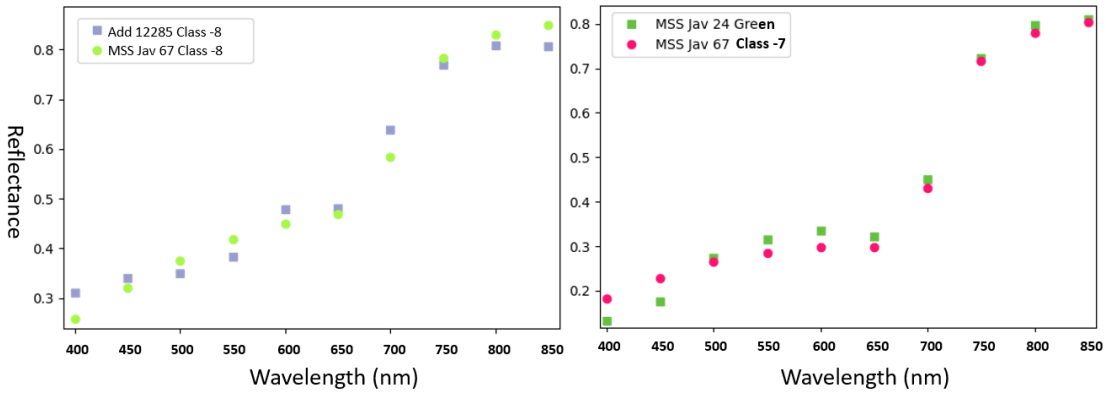

Fig.5.34 Graphs showing the similar reflectance spectra between incorrectly assigned indigo pixels in MSS Jav 67 f12r with lilac pixels from Add MS 12285 (Left) and green pixels found in MSS Jav 24 (Right).

5.4.2.7 Class -7 (Indigo + Orpiment)

Indigo and Orpiment mixtures can be seen over most of the manuscript collection, with correct classification performed by the SVM in almost all cases. Although the class table mainly shows that proper assignment of pixels into the class is performed for most manuscripts, there are some scenarios where some assignments made to the class are instead false positives and are in fact classified with indigo and orpiment mixtures due to them consisting of a mixture of indigo with a different yellow.

One good example of this is in manuscript Or 15932 where many of the green mixtures used consist of a yellow and indigo mix, but the yellow used is not orpiment, and instead is suspected to be an organic pigment. An example of incorrectly classified pixels can be seen in Fig. 5.35, where for folio f21v, the pixels for the green leaves on the outer parts of the tree on the right-hand side of the page are placed into the indigo and orpiment mixture class even though the yellow used is not orpiment. The spectrum for both this Or 15932 green and a green indigo and orpiment mixture extracted from some training data in MSS Jav 24 can be seen plotted together, showing their similar spectra also in Fig. 5.35 and clearly demonstrate that there is a similarity spectrally between the two mixtures.

This misclassification comes as no surprise as there is no information available to the model that would flag the mixture as being an alternative to indigo and orpiment, therefore it would be assigned to the closest material. This could be addressed by having a mixture of indigo and organic yellow in the training dataset but is also another argument in support of using multimodal classification in the future, as clustering also encountered the same problems in chapter 3.

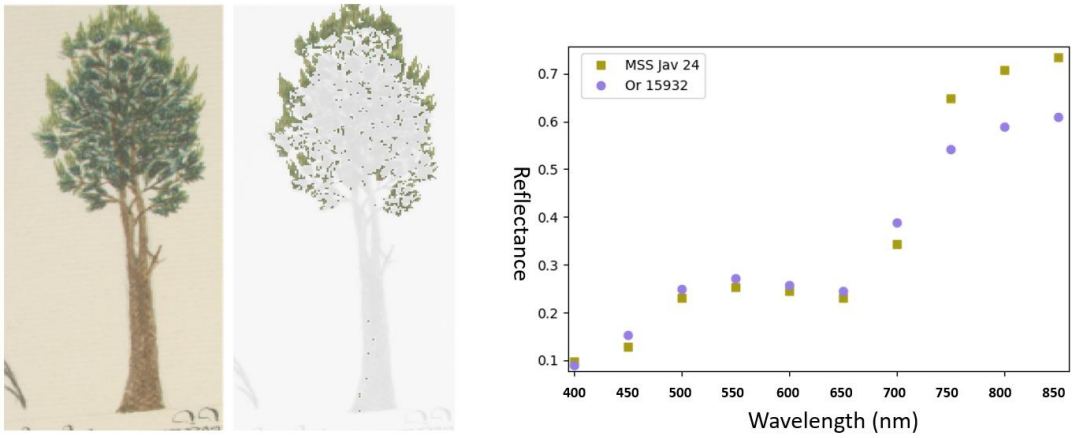

Fig.5.35 Left: Images showing the false positive classification of green pixels in Or 15932 into the indigo and orpiment class due to their similarity in reflectance spectra between themselves and training data seen in MSS Jav 24. Right: A graph illustrating the similarity in reflectance spectra of the tree in Or 15932 and the green training areas in MSS Jav 24.

5.4.2.8 Class -8 (Indigo + Vermillion)

The lilac-coloured mixture of indigo and vermillion does not appear many other times throughout the collection apart from where it exists in the training dataset. However, there are two circumstances where there are groups of pixels assigned to the class. This occurs in Add MS 12339, where a correct classification of the mixture is performed, and MSS Jav 89, where an alternative mixture not identical to that of the training dataset is assigned due to its similarity in spectral features.

In Add MS 12339, the correct classification of the lilac mixture is produced for the long-tailed creature drawn and illuminated on the outside of the double frames on folio f1v. This becomes clear when extracting the spectrum from the lilac training data and comparing it directly with the spectra of the same class in Add MS 12339, where it is highly suggested from the VNIR reflectance that the two pigment mixtures are likely the same due to their similar reflectance spectra, as shown in Fig. 5.36.

However, when comparing training data from the lilac regions with the same class on MSS Jav 89, the two spectra, while similar, likely do not correspond to the same material composition due to the changes in spectral features which are present at both 400nm and 550nm (as shown in Fig. 5.37). Instead, it is likely that these pixels were assigned into the lilac class due to there being a very thin application of indigo, and with no thinner application of indigo being included in the training dataset, it was placed more closely into the class for the lilac-coloured vermillion and indigo mixture, which was applied in thin layers onto the paper in the training manuscript.

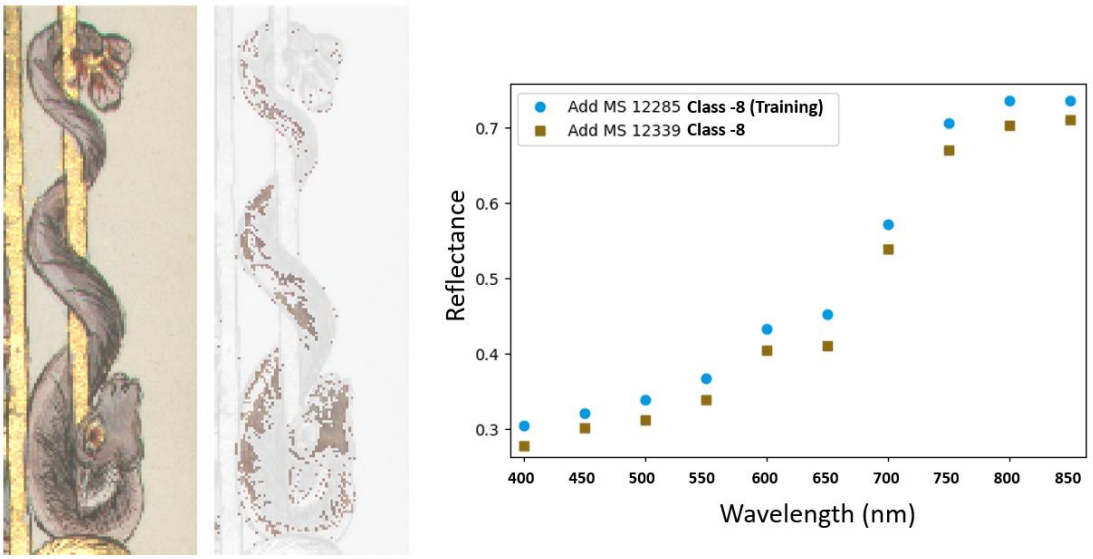

Fig.5.36 Left: images showing the correctly classified lilac regions in Add MS 12339. Right: Graph representing the mean reflectance spectrum of the lilac pixels in Add MS 12339 in comparison to the lilac training data taken from Add MS 12285.

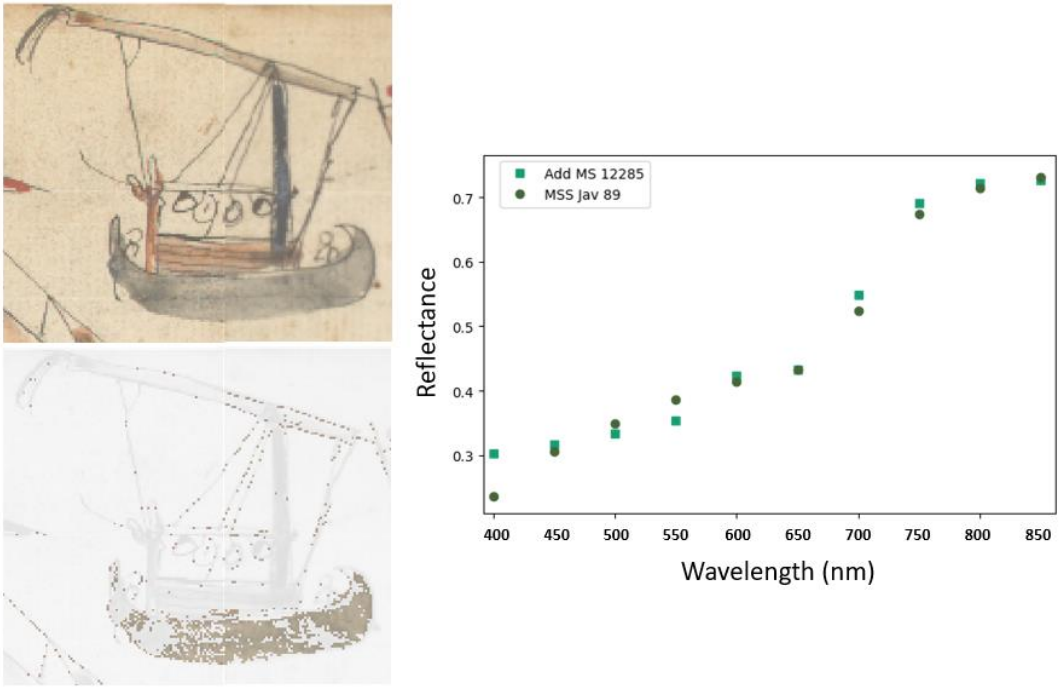

Fig.5.37 Left: Images showing the incorrectly classified indigo pixels in MSS Jav 89. Right: A graph showing the reflectance spectra for the incorrectly assigned indigo pixels with the mean spectrum of the lilac class pixels seen in Add MS 12285.

5.4.2.9 Class -9 (Indigo + Vermillion + Orpiment)

When looking over the entire collection, the reddish-brown pigment mixture can be seen successfully classified in numerous manuscripts. One interesting example can be seen in manuscript Or 15026 f192r, where a correct classification of the reddish-brown mixture is performed for the alternative art style to that seen in the training dataset, as shown in Fig. 5.38.

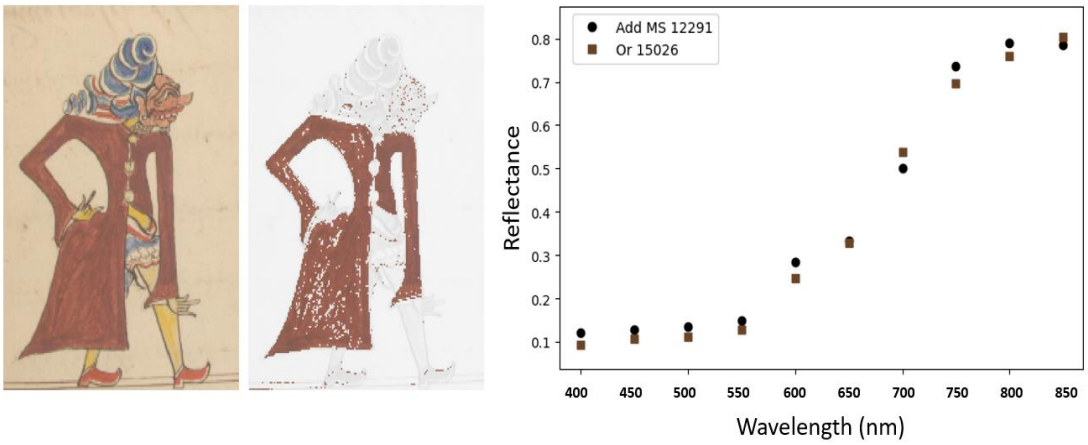

Fig.5.38 Left: Images showing correctly classified reddish-brown pixels in Or 15026 f192r. Right: Graph showing the reflectance spectrum of the reddish-brown pixels in Or 15026 compared to the same class in the training data from Add MS 12291.

As the presence of mercury, arsenic and sulphur can be detected in XRF point measurements taken for this brown coat, there is substantial evidence which supports the reddish-brown classification as being true, which illustrates the strength of using the SVM trained on the grouped clusters.

However, as is commonly seen in different mixtures over the collection, there are also multiple manuscripts which possess primarily false positive assignments instead of true positives, as highlighted by the yellow cells in the class table (Fig. 5.20). In most of these scenarios a spectrum like that of the reddish brown is normally achieved due to the layering of vermillion over indigo, instead of a direct mixture of both with orpiment, as is demonstrated in Fig. 5.39, where this layering can be seen in manuscript Add MS 12379.

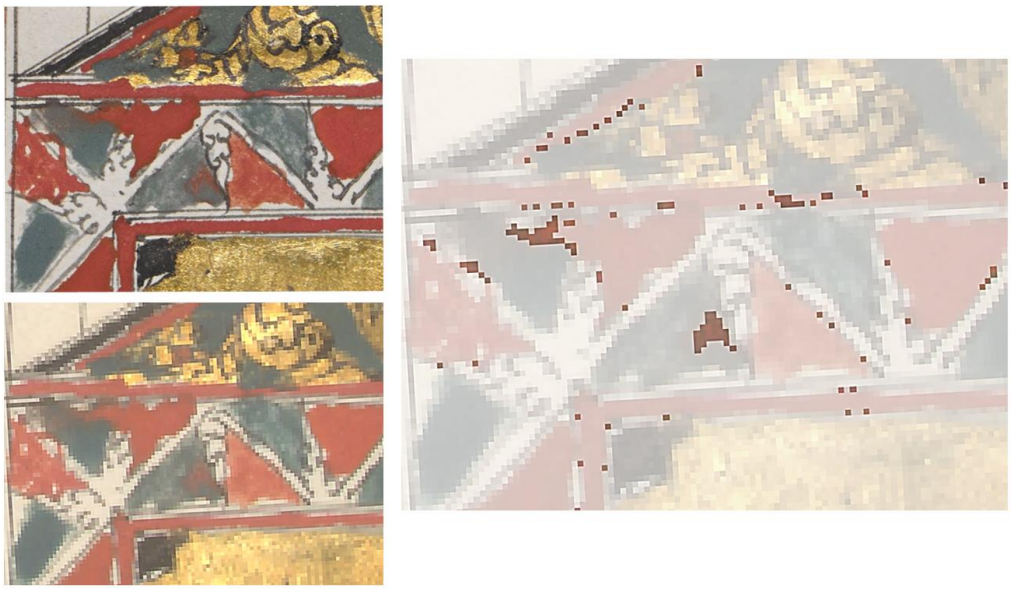

Fig.5.39 Incorrectly classified pixels in Add MS 12379 which have been grouped into the reddish-brown cluster due to the layering of vermillion over indigo.

Interestingly though, even though the appearance of the colour and spectrum indicates a use of indigo, extracting an XRF spectrum from the blue frame elsewhere on the page showed arsenic and sulphur peaks, suggesting that there is some orpiment within the mixture of the blue pigment mixture. If this is the case, then it implies that this may have been classified due to the layering of vermillion over a mixture of indigo and orpiment, meaning that the same materials have been detected and classified appropriately.

5.4.2.10 Class -10 (Indigo + Vermillion + Orpiment (Dark))

For class -10, most of the pixel assignments are false positives which possess similar spectra to those seen in the training data for Add MS 12280 except towards the NIR parts of the spectrum. One good example to illustrate this is in MSS Malay D4 f4r, where a different VNIR spectrum is incorrectly classified into the indigo, vermillion, and orpiment class, as demonstrated by Fig. 5.40. Once again this is most likely an issue being encountered due to a lack of training data. For this manuscript however, a useful part of the classification method became clear, as the use of high probability class maps resulted in few incorrectly classified pixels in MSS Malay D4. This does suggest that the pigment mixture is different for many brown pixels, however there are clearly still other approaches that should be taken in future to alleviate misclassification problems such as this.

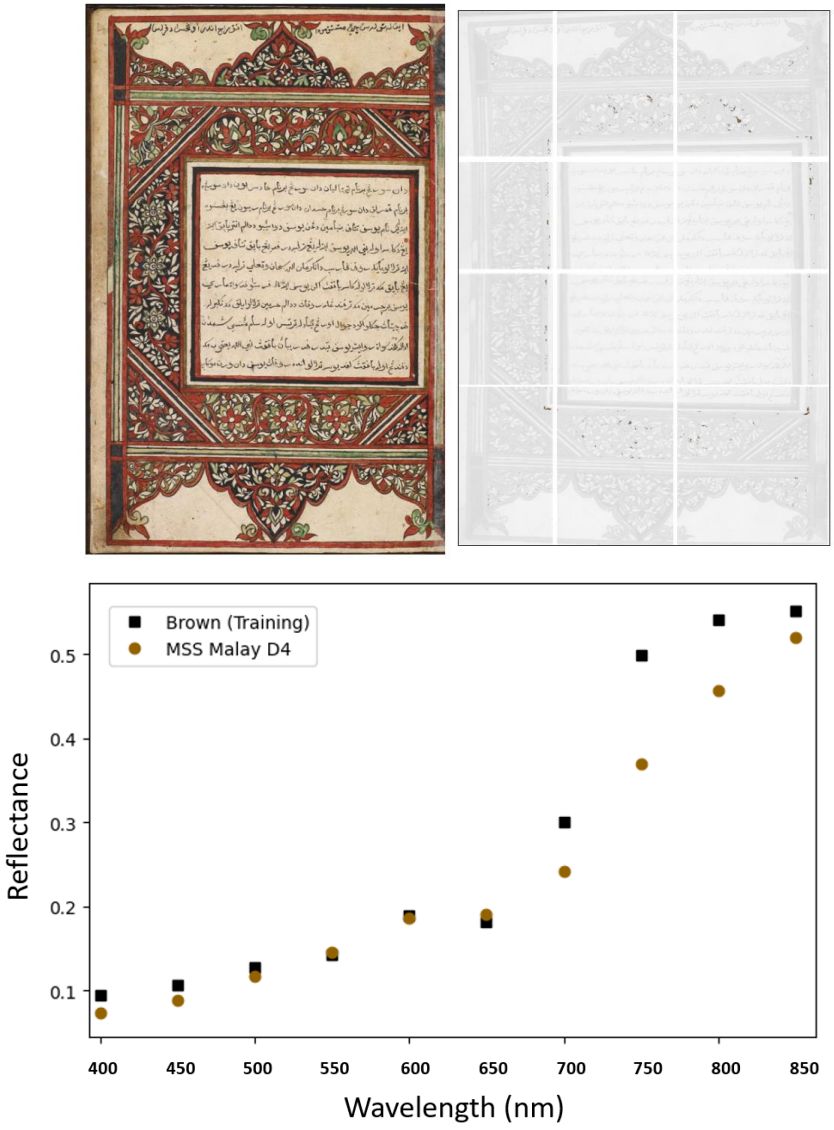

Fig.5.40 Images showing the RGB image (Left) and tru colour class map (right) of incorrectly classified brown pixels in MSS Malay D4 f4r.

5.4.2.11 Class -11 (Bone Ash White)

Throughout the collection, many of the pixels assigned to the white class are not actual white pigments but possess similar enough spectra that they are assigned regardless. This is expected as there is normally a lack of absorption features in white pigments which can be used to find similar spectra, causing other bright featureless spectra to become assigned to the same class. A good example of this has already been shown for the paper in 5.4.2.1. Another example of where this can commonly occur is in the interfaces between black inks and clear paper, where the greyish pixels created due to the change from one colour to the next possess spectra with few features and similar intensities to the white areas seen in the training data.

There are also cases where a different white is used in the collection but is still classified into the same class as the bone ash white, mainly since the white pigments have very few spectral features and sit at roughly the same intensity. For example, in manuscript Or 13295, there are many small white areas painted within the double frames on folios 190v and 191r, in small crosses or within the flower motifs, which are classified into the bone ash white class, as shown in Fig. 5.41.

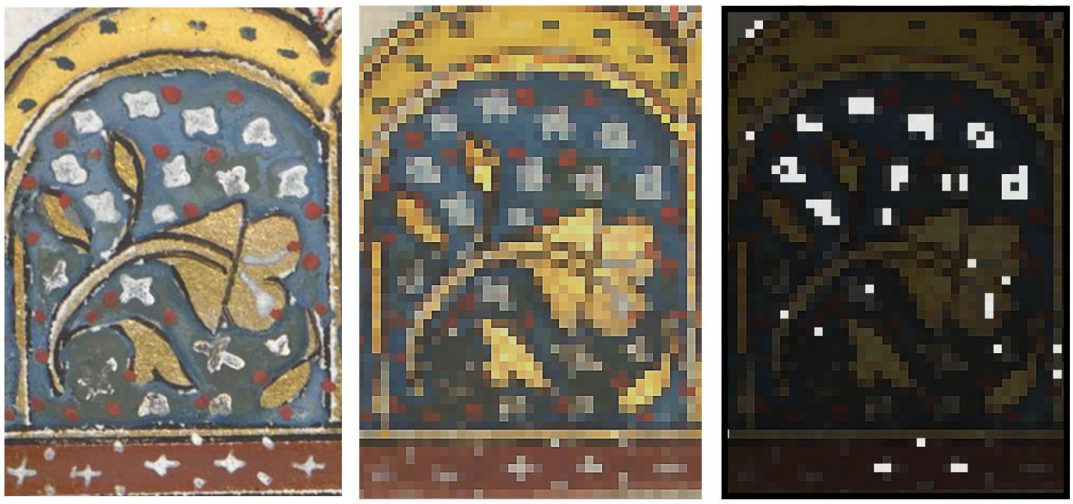

Fig.5.41 Images showing the misclassified lead white parts of Or 13295 which were assigned into the same class as the bone ash white found in the four training Javanese manuscripts

This is considered a misclassification as the detected white pigment in these areas is more likely lead white, as the counts for lead increase in the majority of XRF spectra recorded for this manuscript where white is used to achieve lighter colours. It is expected however that this misclassification would occur, as lead white can be difficult to distinguish from others without the use of other spectroscopy techniques.

5.4.2.12 Class -12 (Gold)

As gold is present throughout most of the collection, it comes as no surprise that there are many manuscripts which have pixels assigned into the gold class. The class table illustrates the successful classification of gold clearly, but in many of the manuscripts it also notes that there are additional false positives that can be included alongside the true positivise class assignments. For the manuscripts which are labelled with having "mostly false positives" in the class table, the cause of the misclassifications tends to be a result of yellow pigments being assigned into the gold class, which is to be expected considering that many yellow pigments and gold, under illumination that avoids any specular reflection, can result in very similar VNIR reflectance spectra (as demonstrated in Fig. 5.42).

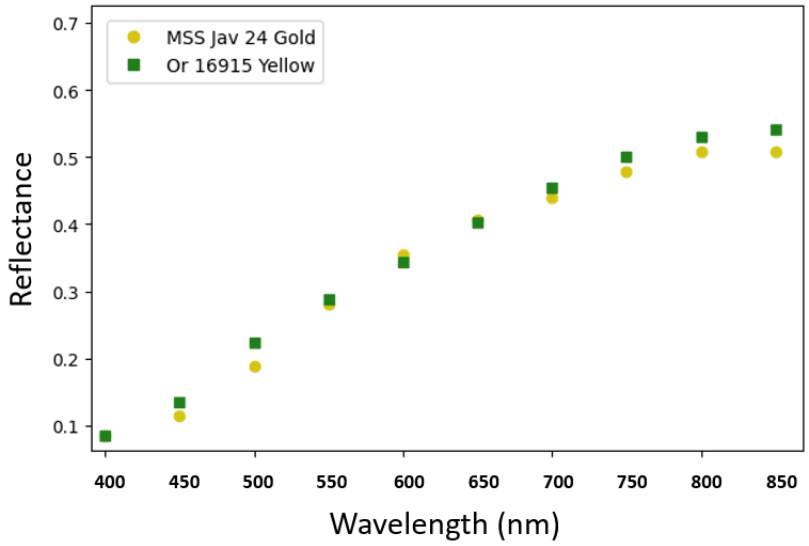

Fig.5.42 Graph showing the reflectance spectrum of gold in MSS Jav 24 and a similar yellow colour in Or 16915 which was incorrectly assigned into the gold class.

5.4.2.13 Class -13 (Vermillion + Black)

The dark red class appears in almost every single manuscript where vermillion can also be found, and normally exists as a consequence of degradation, or black and red being placed either next to each other or atop one another in different layers, a good example of these classifications can be seen in MSS Malay D4 (see Fig. 5.43). There are very few times where a dark red appears to be used as an intentionally mixed pigment, however there are a few circumstances where this is the case outside of the training data examples. However, alongside these correct classifications, there are also some misclassified regions in manuscripts which come as a result of regions sharing similar reflectance spectra. A good example of this can be seen in Or 13295 f190v, where the darker brownish-red parts of the frame are classified into the dark red class due to their VNIR spectrum fitting well with the dark red in the training dataset. It is not a surprise that these areas

may be classified into the dark red, as the mixture was confirmed in chapter 3 to be a combination of Prussian blue, orpiment, and vermillion, which was not implemented into the training. However, it is unfortunate that the implementation of probability did not label the brown-red mixture as unclassified. The class maps for this manuscript page can be seen in Fig. 5.44 and shows a consistent classification over much of the manuscript page which is split into numerous data cubes.

Fig.5.43 Images showing (Left) an RGB image of MSS Malay D4 with correctly classified dark vermillion pixels in a true colour class map (right).

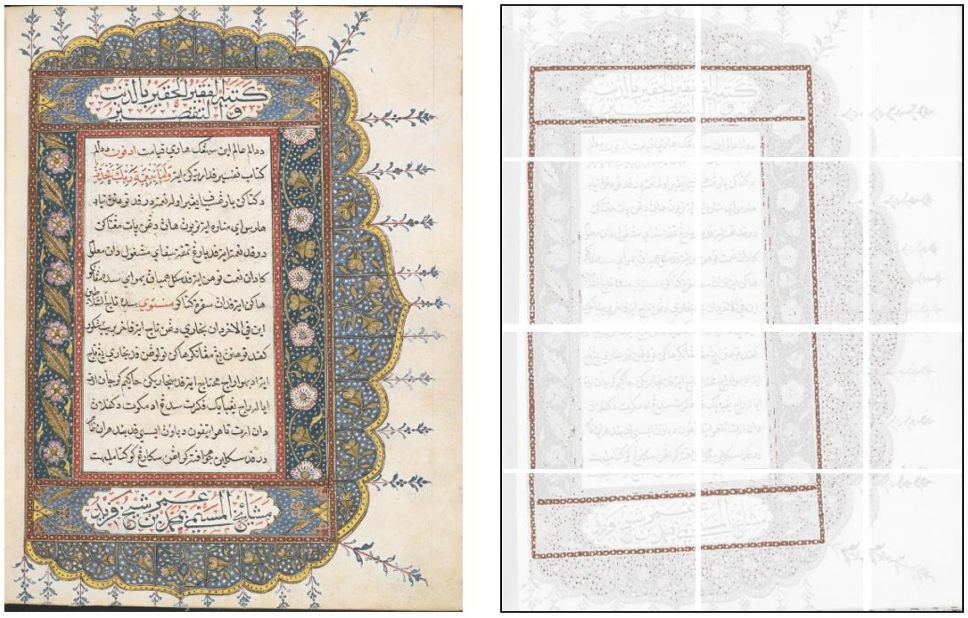

Fig.5.44 Image showing (left) RGB image of Or 13295 f190v alongside a true colour class map (right) of the incorrectly classified vermillion, orpiment, and Prussian blue mixture.

5.4.2.14 Class -14 (Black Ink/Pigments & Background)

As the final class acted to encapsulate the majority of dark areas, the performance of the SVM is not particularly informative. However, there are a few areas in which some of the pigments with higher absorption, such as dark Prussian blue determined during the pigment analysis in chapter 4, are grouped into class -14 due to them possessing similarly low reflectance intensity and consequently weak spectral features. As it has already been covered in detail throughout the previous sections, Or 13295 will be used as a good example yet again, where throughout the double frames, Prussian blue is the main pigment used and is classified with the other black materials and the background, as shown in Fig. 5.44.

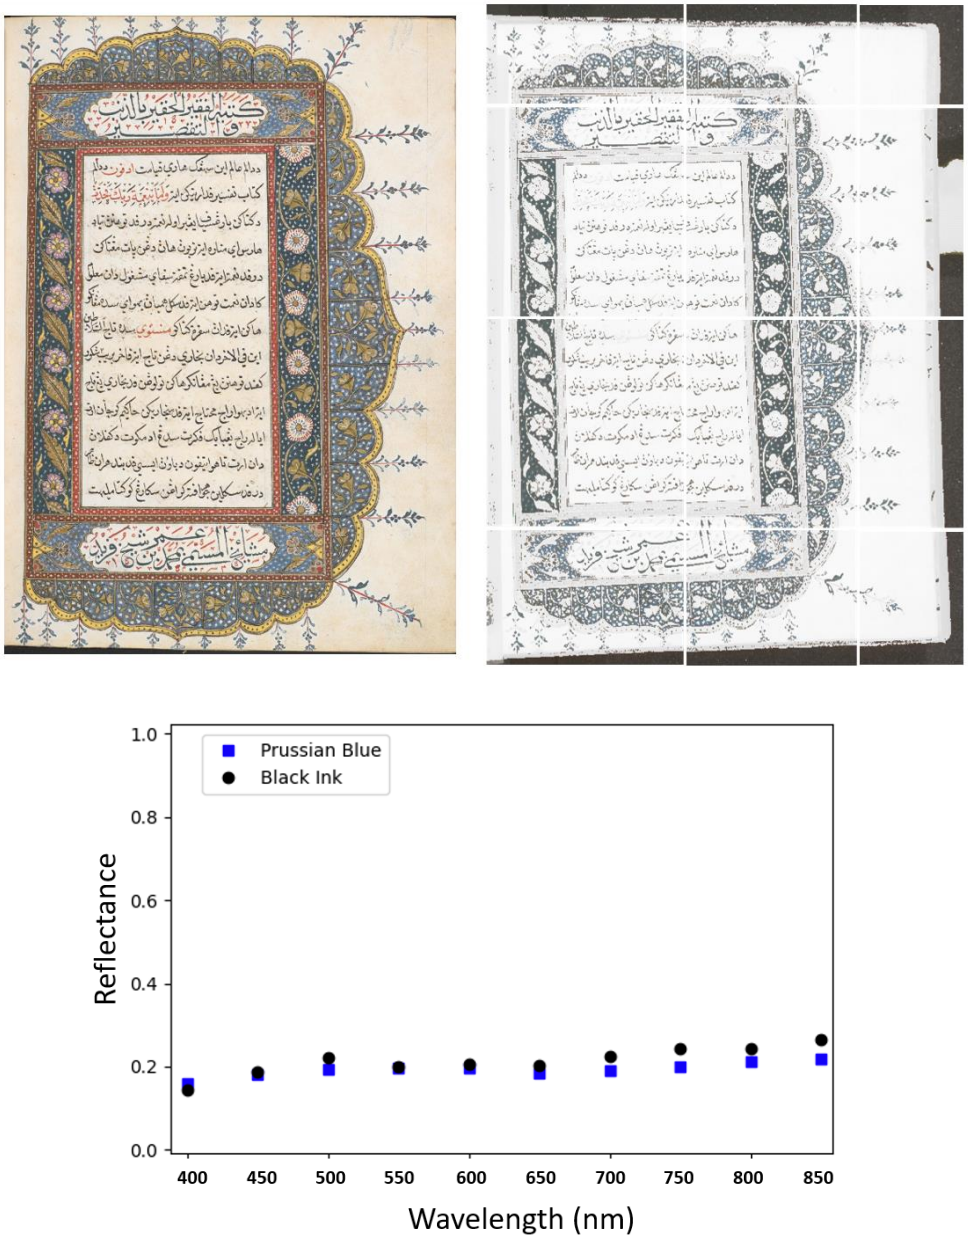

Fig.5.45 Top: Images showing (Left) the RGB of Or 13295 f190v alongside (right) the true colour class map of the Prussian blue regions which are incorrectly classified into the black due to their high absorption. Bottom: A graph showing a black and Prussian blue area plotted together with similar highly absorbing reflectance spectra.

5.5 Conclusions & Future Developments

This chapter explored the use of classification techniques on VNIR spectral imaging data recorded for a large manuscript collection, and how they may fit into the workflow for performing large-scale pigment mixture identification in the future. In the first half of this study, different classification techniques were compared to evaluate their performance for identifying pigment mixtures in illuminated manuscripts, where it was found that four different techniques: k-Nearest Neighbours, Random Forest, Multilayer Perceptron and C-Support Vector machine, were all found to perform well. As the highest accuracy achieved in this evaluation was the C-SVM, though only by 0.1%, it was decided to use this as a model with which to perform classification for 14 different classes, created through clustering and grouping 4 Javanese manuscripts, on the large manuscript collection already explored in chapters 2 and 3. The results from this test showed that C-SVM based classification can perform well in classifying unseen VNIR spectral data, implying that C-SVM or the other highly accurate models, may cooperate well with the methodology already outlined in the previous chapters. Overall, it seems that classification techniques in use with large-scale clustering of collections can be incredibly useful in future studies of illuminated manuscripts. However, care must be taken when using classification to perform material identification over many manuscripts at once, as there is no guarantee that the performance will always be accurate if training data is limited for use on new datasets, even in the case where high probability class maps are used, as was demonstrated for black ink/pigments and the background.

During this chapter it was discussed that while the C-SVM performed well, it was also both time and computationally expensive for performing classification at a large scale, even after reducing the number of spectra in each image by a factor of 16. While this is a drawback of performing classification as opposed to clustering, the problem can likely be solved by performing classification on clustering results. While such an approach would save computational effort and time, it also presents an opportunity to link classification together with the methodology developed for clustering and grouping, where instead of simply using clustering and grouping results to train a model, the two approaches could work in unity together for future for large-scale studies. The intention would be to expand the clustering and grouping workflows, such as the Kubelka-Munk (KM) informed grouping workflow covered in section 2.6.4 of this thesis, to include a classification step. The outcome of this unification of both techniques would therefore be a looped process in which a model, such as C-SVM, learns from pixels within manually grouped clusters, allowing for new future clusters (not pixels) to be automatically classified in a time efficient way. Successful classification will therefore be able to easily group data over entire collection quickly, but in addition to this, clusters which cannot be classified due to possessing low probabilities can also be returned to the user to undergo a new round of KM informed grouping as per 2.6.4. This would in turn allow for a retraining of the classification model so that

the next time the same pigment mixture or material is encountered, it is classified correctly, and with a high probability. Each time this process is performed, classification and clustering would work together in a loop, until every spectrum in a new collection has been not only identified, but also stored within a database and trained model for future use. This proposed expansion of the diagram shown in section 2.6.4 can be seen in Fig 5.46 and demonstrates the direction in which this research is likely to continue moving forward.

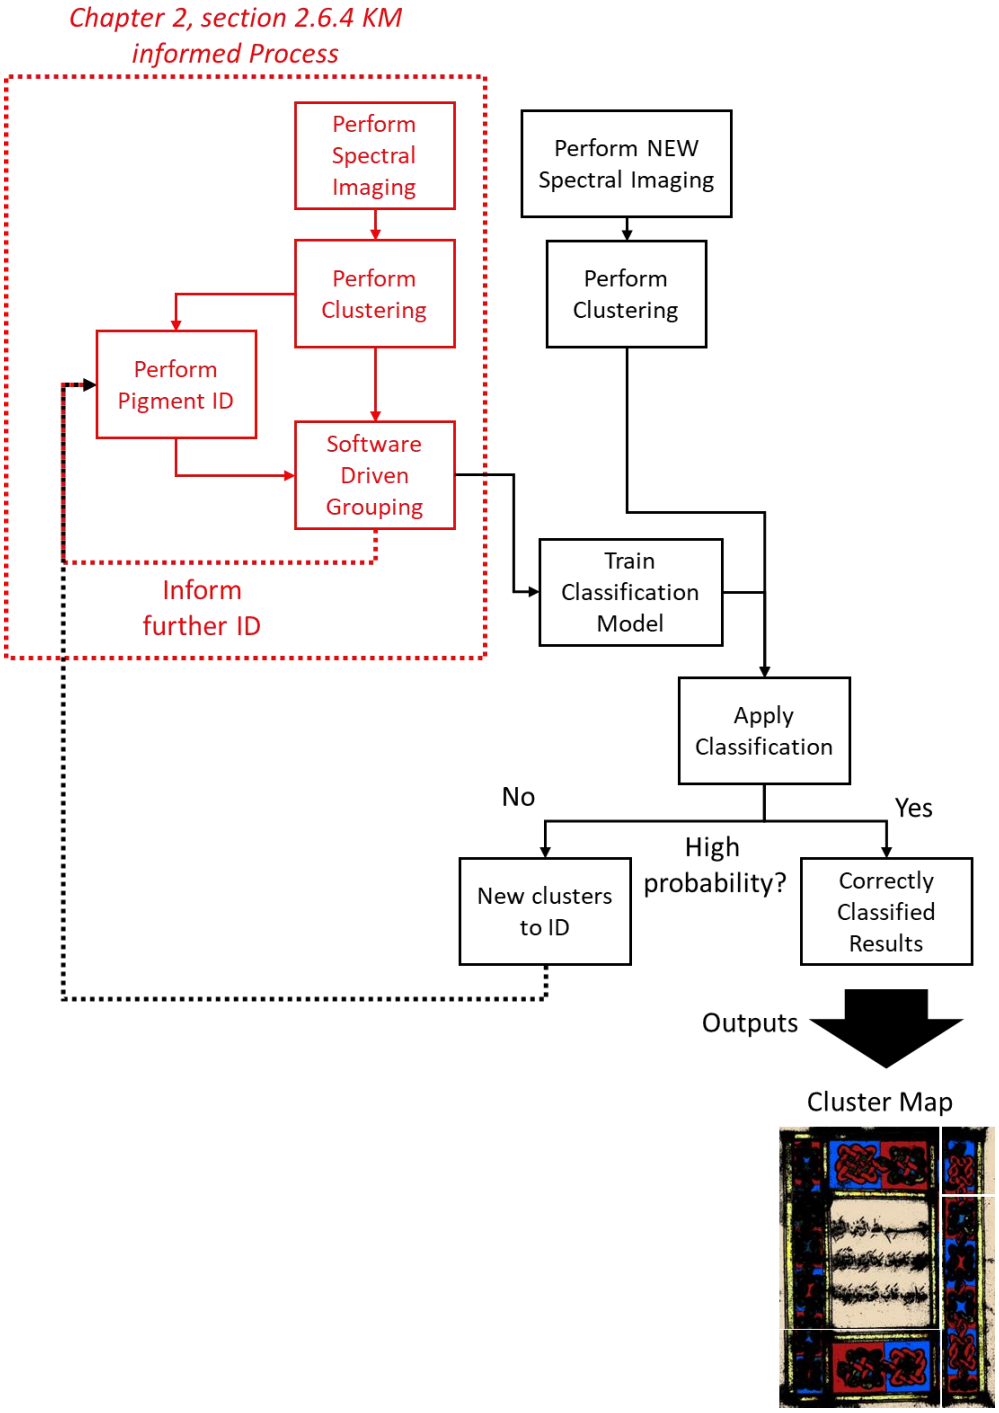

Fig.5.46 Expansion of the clustering and grouping procedure outlined in chapter 2, where Kubelka-Munk informed grouping can inform classification, and vice versa.

Chapter 6
Guisi: A GUI for Analysis of Spectral Imaging & Clustering Data

6.1 Introduction

As large-scale studies of artistic materials used throughout entire collections become more commonplace in the future, there is a growing need for data visualisation tools that can be used alongside clustering or classification techniques to allow for the interrogation of high dimensional and highly complex spectral imaging data.

At the time of writing this thesis, multiple approaches can be taken to analyse spectral imaging data in cultural heritage applications using a variety of different software, including those designed for different specific techniques and instruments such as ESPRIT Spectrum (Bruker, n.d.), PyMCA (Solé et al., 2007), Specim IQ Studio (Picollo et al., 2020; Specim, n.d.) and more. Furthermorea variety of software has also been developed by the research community on an ad hoc basis using common software platforms such as MATLAB, Python or R. Many of these software programs can offer the ability to perform data processing, extraction of spectra, spectral image visualisation and even allow for users to perform clustering or classification on data directly (such as Prediktera's breeze and EVINCE software (Prediktera, n.d.), or perClass BV's MIRA (perClass BV, n.d.)), however, there are seemingly no clear leading software applications that can allow for fully accessible analysis of any spectral imaging modality and any type of spectral image based clustering result produced from any algorithm. Furthermore, when considering the use of software for the analysis of large-scale datasets, these applications cannot facilitate and allow for simultaneous spectral imaging or clustering visualisation and analysis over many items at once, which was essential for most of the work in this thesis.

As a large part of this thesis scope was to design a working methodology for analysing multimodal spectral imaging data and large collections, many analytical tools were investigated and newly developed or redeveloped throughout the different projects mentioned in earlier chapters. Eventually, compartmentalising many different small codes or programs for different analytical purposes became cumbersome, and a streamlined approach to go from data collection to material analysis was desired. As a result, the design and development of a new software application, named Guisi (GUI for the Analysis of Spectral Imaging & Clustering Data), was prescribed as the best route forward in the analysis of large-scale datasets. Guisi was initially created as an application that would allow for the basic visualisation and analysis of spectral imaging data through simple isolation of colour information, spectral bands, and extraction of spectra.

However, as it progressed the software became essential to performing clustering analysis, especially when performing processes such as the grouping or plotting of clusters for understanding material distribution. As the software became central to the methodology of the different projects, it is covered further within this chapter, where its uses and design choices are discussed for the current version "Guisi v.0.3".

As this chapter is mainly revolved around the implementation of newly developed software central to the clustering approach outlined in chapter 2, the structure of the chapter is slightly different to those seen previously. Sections of this chapter will discuss the design choices and implementation of Guisi and will also include discussions on how Guisi may be used in the broader cultural heritage community. Arguments in favour of transitioning to more universal data types will be made so that analysis can be more standardised across joint projects in the future, aiding workflow efficiency across institutions. A final discussion about future developments will also then be included with the intentions of how the software intends to be used in the future.

6.2 Data Pipeline, Needs & Design Choices

When considering the design of software for the analysis of spectral imaging and clustering, it was first important to understand what the different needs and requirements were for multimodal and large-scale studies so that the software could be optimally designed to fill a specific role.

6.2.1 Acquisition to Results Workflow (ARW) for Spectral Imaging Studies

In the grand scheme of performing different spectral imaging studies, several stages were always followed so that data could be captured, processed, and then eventually used for analysis within the research encountered in this thesis. Throughout this chapter, these different processes will be referred to as the 'Acquisition to Results workflow' or ARW. Covering the entirety of a spectral imaging study, the ARW normally begins with data collection (spectral imaging), and in the context of this thesis typically ends with clustering, grouping, and material identification (pigment identification). A diagram illustrating the typical ARW for this thesis can be seen in Fig. 6.1.

At every step in this ARW, it was found that both data visualisation and interrogation of raw or processed spectral data was useful. For example, aside from being beneficial for analysis, it was often also useful to understand if proper calibration, data collection and processing were taking place while in-situ so that when clustering and material identification is performed later, it was done so with data which had been captured and processed in the most optimal way. Because of this, it was imperative that Guisi could be used with equal accessibility over the entirety of this ARW for different techniques, therefore implying that the software and the data for clustering had to be designed and used in such a way that allowed for any spectral imaging type to be used in

the software, from XRF to Vis-NIR. While Guisi's uses are covered in more detail in sections 6.3 and 6.4, the first investigation covered in this chapter discusses how this accessibility to data at any point of the ARW be achieved by ensuring that a universal datatype is chosen to be used with Guisi, so that spectral imaging data of any format can be used in the application at any stage.

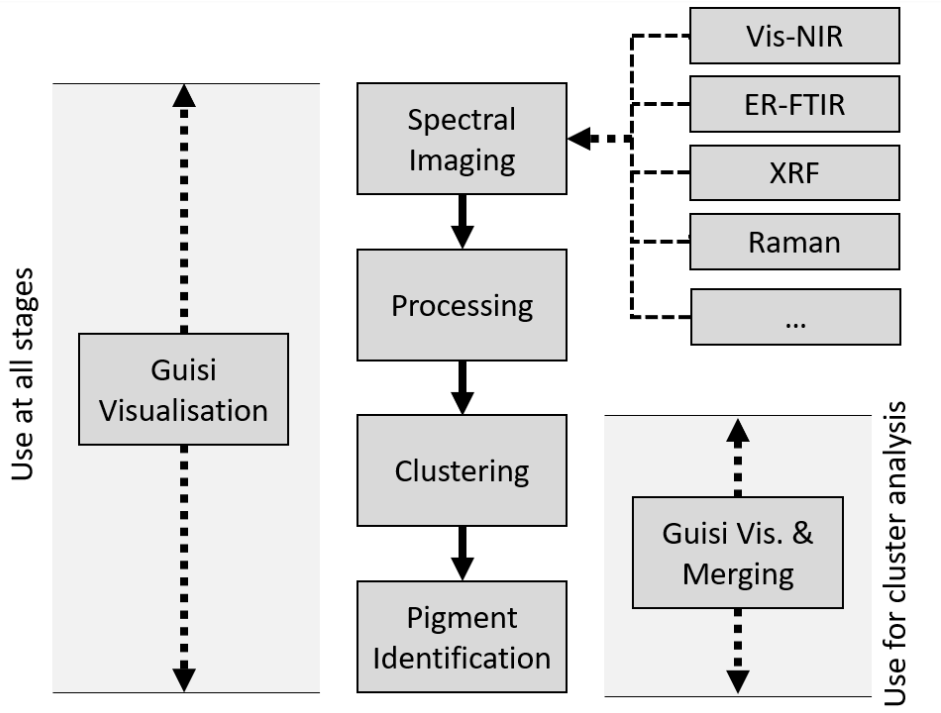

Fig.6.1 Diagram representing an Acquisition to Results Workflow (ARW) for spectral imaging studies of large collections or holistic multimodal data

6.2.2 Data Formats for Holistic Spectral Imaging Studies in Cultural Heritage

Though most spectral imaging data for collections can essentially be simplified into a set of volumetric data cubes, one of the greatest obstacles in allowing for a universal file format to be used for cross-system data analysis in Guisi is how to account for the many different formats already commonly implemented in different applications. Across different disciplines and over many techniques in cultural heritage, spectral imaging data is stored and saved in a variety of different formats, with many techniques resulting in data which can require different pieces of software to be opened, investigated, and analysed (some examples of which are shown in Fig. 6.2). Many of the software packages required to interpret the data can be tied into the sale of specific instruments or are specifically designed to be used with only a single type of spectral imaging data or for specific purposes. This segregation of data types and tools is understandable when considering that single-mode analysis is used often in many studies, however for large-scale holistic spectral imaging studies in cultural heritage, it presents problems where data cannot be easily used together as it is always compartmentalised in some way. Even within this thesis alone

there were numerous different data formats which were tied to various instruments or techniques, resulting in holistic analysis and multimodal clustering of spectral imaging data becoming needlessly complex as it required consistent conversion between data types so that different analysis processes could be performed efficiently.

Instrument	Technique	Spectral Imaging File Format
Bruker M6 “Jetstream”	XRF	ENVI
Bruker CRONO	XRF	HDF5
PRISMS	Vis-NIR	VIPS
Specim IQ	Vis-NIR	ENVI / RAW
Bruker Lumos II	FTIR	ENVI / OPUS
MegaVision EV	UV-Vis-NIR	DNG

Fig.6.2 Table demonstrating some of the different spectral imaging file formats produced from different techniques in cultural heritage.

In addition to the concerns regarding bringing data together in a complementary way, there are other difficulties to encounter when sharing data of different formats for joint analysis efforts between collaborators or coworking researchers. If data is stored in many different formats before being transferred to coworking researchers or groups, it may have to be repeatedly converted into accessible data types if associated software to the instruments used is not universally available. As the conversion between different large datasets can be computationally and time expensive, it is therefore important for common formats to be used across institutions and research groups to increase productivity for large collections. Therefore, to avoid processing and analysis hurdles like this in the future, it makes sense to use only a single universal data format for these types of studies. However, to understand which formats may best work for an application such as Guisi, some of the strengths and weaknesses of different data types encountered during the research performed in this thesis were first investigated to understand which would best serve the ARW moving forward.

6.2.2.1 ENVI

The ENVI file format is often considered the industry standard for most spectral imaging applications and is therefore a decent candidate for standardising the storage of spectral imaging data. However, ENVI is still not used in all scenarios, with many different commercial or lab-built instruments opting for different data formats completely, as illustrated previously within the table in Fig. 6.2. ENVI can be powerful in storing spectral imaging data of different formats as it is completely lossless and has plenty of options for storing different sizes of raw data captured from a variety of different instruments. However, these features can also present disadvantages

when used in practice by a researcher who has not necessarily collected the spectral imaging data themselves, as accessing ENVI data always requires unique inputs catered to the data collection settings to open saved data correctly. When performing data analysis for large collections across multiple different spectral imaging data types this can begin to overcomplicate the ARW, as many different properties must be accounted for including the samples (columns), lines (rows), bands (channels), byte orders, byte offsets, interleaving configurations, data types (bit depth) and type of header file extension (*.hdr or *.rpl). While Guisi could use the ENVI file format for certain applications, a detailed understanding of where the data is sourced and how it was collected must always exist, so that a user can input the correct parameters to acquire the correctly formatted spectral imaging data. For large collections this can prove to be time-consuming, as changing parameters for hundreds or even thousands of spectral imaging data cubes would not be ideal.

6.2.2.2 HDF (Hierarchical Data Format)

Another widely used format, HDF shows promise for analysing multiple types of data as it was designed specifically for the storage of scientific data across different systems and techniques. It can be used to store volumetric data (or arrays), tables, annotations, raster images and more, and therefore can be used as a single file to hold all the information associated with a spectral imaging data cube. Additionally, HDF is also widely supported, with many prebuilt functions in many programming languages able to access the data easily, including R (Fischer B et al., 2022), MATLAB (MATLAB, n.d.), python (h5py, n.d.)and more. The main disadvantage of using HDF in the context of this study however is that the structure of the file itself can vary depending on the system used and how the file is created. This means that without ensuring that all HDF file structures are the same for different independent spectral imaging techniques, an understanding of where certain data exists within any HDF file is critical, which can make automatic accessing of data problematic for large-scale holistic studies. However, as programs can be written to automatically read the tags and structure of an HDF file, it does hold the potential to be a universal format in the future if a standard conventional file structure is followed in cultural heritage studies.

6.2.2.3 TIFF (Tagged Image File Format)

An alternative file format widely used for storing imaging data is the TIFF file format, a format which can store volumetric data and any tagged data, such as header information, as a single '*.tif' or '*.tiff' file, meaning that any spectral imaging data and its associated reading properties, wavelength or frequency channels can be kept together at all times, a seemingly minor detail that is increasingly important when dealing with large sets of multidimensional spectral data cubes. The traditional downsides of the TIFF file were that it uses 32-bit offsets, which means that a TIFF cannot store information with a bit depth greater than 32 bits (float 32, uint32, int32) and is

limited to a maximum file size of 4GB. However, as big data has become more common in the 21st century, the development of BigTIFF began to rise as a more common variant of the TIFF file which instead uses 64-bit offsets, meaning that the maximum size of a single BigTIFF file can reach 18 exabytes (equivalent to $\approx 1x10^9$ GB), which is much greater than any spectral imaging dataset which would ever be kept as a single file. The capability of being able to store any spectral imaging data cube in a single file with all it is associated wavelength/frequency information can massively simplify the process of data analysis and streamline data workflows. It also allows for easier transfer of files between different instruments, applications, computers, researchers, or institutions without the need for conversion between the many different data types commonly used in many different sectors.

6.2.2.4 VIPS

The file format originally used to store raw and processed PRISMS (Liang et al., 2014) spectral imaging data used throughout this thesis, VIPS (Cupitt & Martinez, 1996; Martinez & Cupitt, 2005) is an image format that can be used to store volumetric data for easy use with the VIPS image processing library and user interfaces such as nip2. VIPS has seen use in many museums and galleries over the world and is popular due to its open source and cross-platform capabilities, and therefore presents a unique data format which may be effective in its use as a universal format in the future. The disadvantages of VIPS however are that the format is not as widely used as other file types, such as ENVI or tiff, and as such would have limited support for many instruments.

6.2.2.5 Other

In addition to the previously mentioned commonly used file formats, there are a multitude of other file formats which could act as a universal data format for analysing large-scale spectral imaging data. However, many of these were not tried or tested during the research carried out in this thesis so are not covered in greater detail. Some of these other formats can include FITS, NetCDF, DNG, IMG, DF3, just to name a few, with many more 3D volumetric data formats developed for many different purposes.

6.2.3 Guisi Compatibility with TIFF, VIPS, and other formats

Due to the ease and accessibility of using BigTIFF over other techniques in supplying spectral imaging data cubes as a single easy-to-access file, the most stable data format to be used with the

Guisi software is TIFF, and when used allows for the visualisation of spectral imaging data of any type and any wavelength or frequency. To support this, a suite of different conversion functions has also been developed to allow for seamless reformatting between common datatypes and BigTIFF, many of which are implemented into the backend of Guisi directly, but it is recommended that any volumetric data cubes imported into Guisi use either the TIFF or BigTIFF formats to avoid import errors. For scenarios where computer RAM becomes a limitation in performing data analysis, such as for large-scale XRF data, the VIPS file format can also be used to access spectral imaging data within Guisi as well, which can sometimes offer memory performance improvements. If a TIFF or VIPS file is not tagged with any information during import, wavelength or frequency data can also be accessed for a data cube file by using an identically named '*.iwvl' file, a '*.txt' file where each row is a float value corresponding to each channel or image band.

6.3 Software Application & Implementation

Throughout this thesis, spectral imaging case studies were performed which required the interrogation, analysis, and grouping of spectral data over multiple export paintings and many manuscript illuminations. In all these studies, investigations into the use of pigments and or painting techniques were performed much more efficiently using Guisi as it provided a platform for cluster analysis and high-dimensional spectral data analysis. The tools used to perform these studies were purposefully developed with the clustering and spectral imaging of collections in mind and are therefore designed to allow for convenient large-scale analysis. This section will cover the structure of the software, including these different tools, and will provide examples of their use in previous chapters where it will then become clearer as to why a visualisation approach was important for this study and will continue to be in the future.

6.3.1 Graphical User Interface

The front-end GUI side of the application is primarily developed in python and its current state consists of three main interactive windows: a main display, a data visualisation display, and a cluster analysis display, of which the latter two are accessible in different tabs. There are additional tools in the menu toolbar at the top of the window, but as these are simple tools used only to upload data or save/open projects, they will not be covered in detail. An example of the main window and tabs before uploading data can be seen in Fig. 6.3, with different sections of the GUI labelled to appropriately match the following sections 6.3.1.1 – 6.3.1.5.

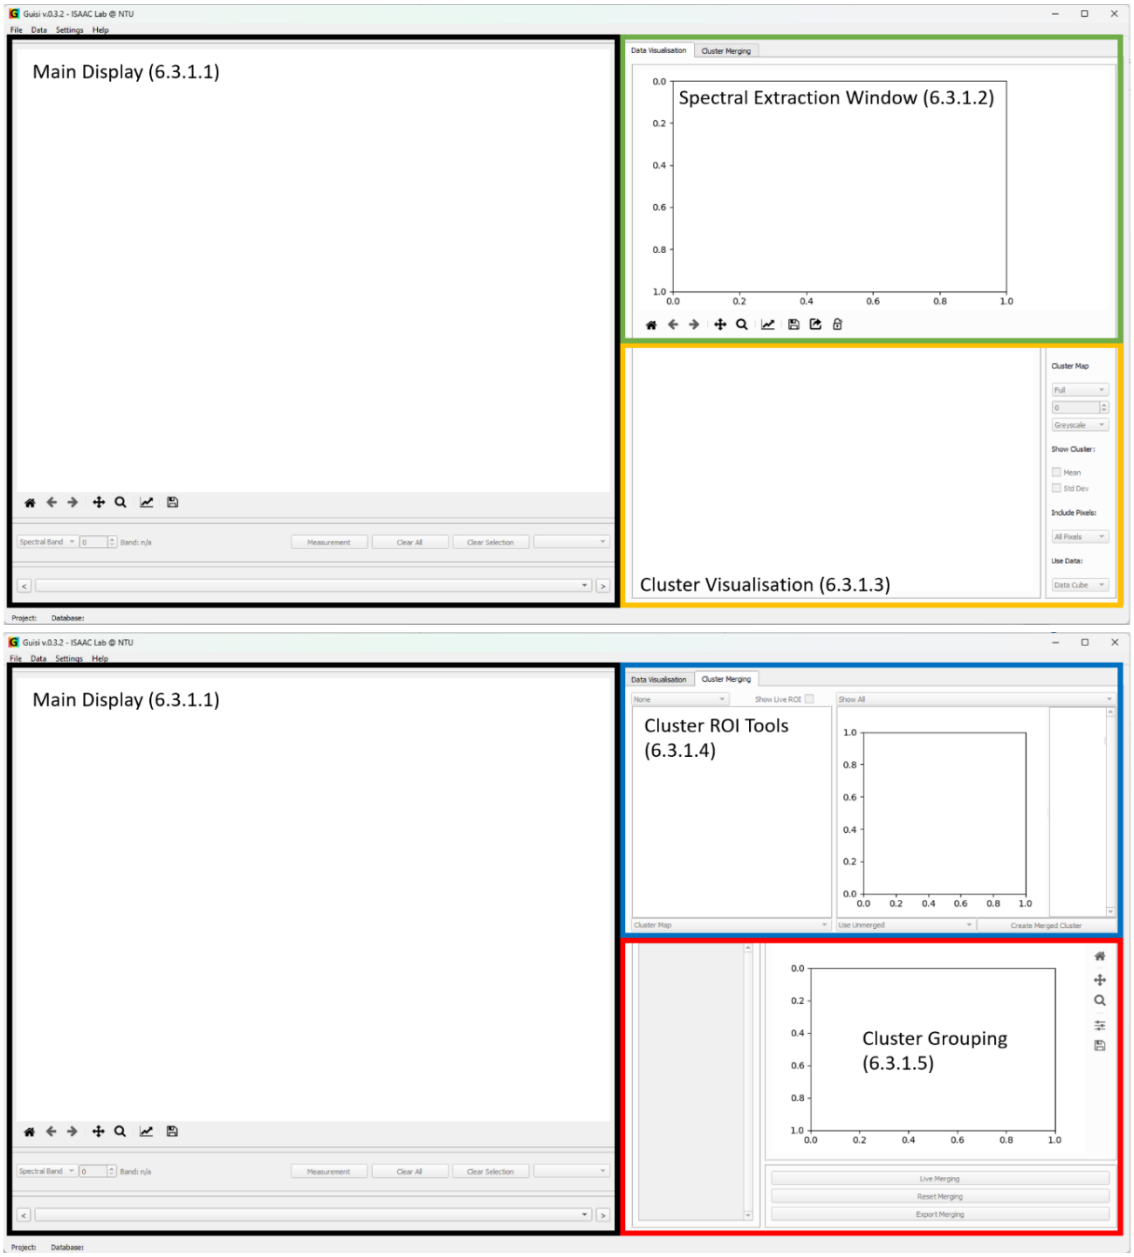

Fig.6.3 images showing the Guisi main window, separated into different sections by colour and associated with chapter sections.

6.3.1.1 Main Display

Highlighted in black within Fig. 6.3, the main display works as both the main data visualisation window but also serves as the primary point of the user's interactivity with any uploaded data. There are multiple tools for selecting images, visualising data of different types, or even extracting spectra to perform analysis directly within the software. From top to bottom, the main display can be split into three main parts: image display, analysis tools, and spectral image selection, all of which can be used for the analysis of spectral imaging data in different ways.

Image Display

Within the main display of Guisi, different types of uploaded data can be visualised in multiple ways. Furthermore, so long as the spectral image files used possess the same names or order, colour images, mean images, spectral image bands, and cluster maps all corresponding to one another can be visualised simultaneously within the main display window, allowing for multiple analytical processes to be performed where:

- Colour images can be used to monitor the basic colour palette used over a painting and pinpoint regions for closer study.
- Mean images can be used to visualise an entire image from data with more sparse information, such as XRF (where in many individual energy channels no signal is found for certain regions, or perhaps even the whole painting).
- Spectral bands can be used to observe individual channels, where singular XRF energy channels or single vis-NIR wavelength bands could be displayed (as was performed for the detection of iron gall inks in chapter 3).
- Cluster maps can be used to visualise the distribution of unique spectra, either as binary maps or in true colour, over an entire image and across multiple spectral data cubes (used throughout the thesis).

In addition to the analytical processes used for material identification after data collection, these features made it possible to check spectral data cubes and individual raw channels simultaneously to data capture, which meant that Guisi could be used as a tool for calibration purposes or for checking and identifying artefacts, saturation, or other problems. Furthermore, certain features such as underdrawings, which do not always require processing to visualise, could also be identified live. Mean images were also useful in quickly visualising XRF data, as the alternative would be to look through sparse XRF energy channels. And finally, cluster maps and spectral image visualisation displayed in Guisi could be used to guide the subsequent collection of point measurements.

Analysis Tools

In the centre of the main display are tools that allow for the extraction of spectra from any type of spectral imaging data. By selecting a region of interest (ROI) when displaying any data in the main display widget, the voxels selected within the ROI will be averaged into a mean spectrum to be displayed in the data visualisation tab. Typically, a rectangular selection of pixels is used to create the mean, however, a user can also capture only the spectra belonging to a specific cluster in one area, allowing for the analysis of clustering results and providing high SNR spectra for only one type of material in an ROI. ROIs are stored in the software and can be labelled or given any colour decided by the user, they can be deleted at will or saved and used for other actions such as to perform cluster merging. This section is also where the images being shown in the main

display are controlled. The tools developed for use here were central to performing pigment identification in the investigations seen in this thesis, as vis-NIR spectra could be easily obtained from areas of PRISMS or XRF data where complementary point measurements using other techniques were also recorded. Furthermore, swapping between multimodal data allowed for the extraction of spectra of different types from the same pixels (such as XRF, Vis-NIR, and FTIR), which was often central to performing cluster analysis during chapter 4. Spectral extraction could be performed for specific clusters in any given ROI, providing high SNR spectra belonging to inhomogeneous areas. This was particularly useful for scenarios where data had to be collected for spatially small objects (only a few pixels) found throughout different manuscript pages in chapter 3 which had a non-uniform structure but repeated as a pattern in many locations. Additionally, this was used consistently throughout the clustering analysis to ensure that the parameters used within the clustering algorithm provided clusters which were accurate and avoided misclassifications.

Image Selection

As the data upload routines for use with BigTIFF *.tiff and VIPS *.v files within the software are completely automated, data can be easily selected and swapped between by using the tools at the bottom of the main display window, allowing for an almost seamless transition between different spectral data cubes and any associated colour images or cluster maps. This streamlines any data analysis as the process of swapping between many different spectral imaging data sets of different types can be made completely effortless. Furthermore, as folders selected for data visualisation in Guisi can be updated live during data collection, spectral imaging data cubes can also be interrogated live, allowing for the analysis of raw or processed data directly during or after collection.

6.3.1.2 Spectral Extraction

The spectral extraction window works as a basic graphical display which can plot extracted spectra (extracted using an ROI in the main display) or mean cluster spectra and their standard deviation ranges, either calculated from the current spectral image alone, or over the entire collection. For ease, the plot seen in this window for any specific ROI selected in the main display will always share the same colour, line thickness, and line style, with the mean spectrum and standard deviation always depicted in black and grey respectively. This allows for a clear separation between the ROI spectra and the cluster spectra if both must be analysed simultaneously, which was often implemented throughout this thesis to validate clustering results and perform grouping. An example of an ROI selected in the main display, next to the extracted PRISMS vis-NIR spectrum for a red region, can be seen in Fig. 6.4.

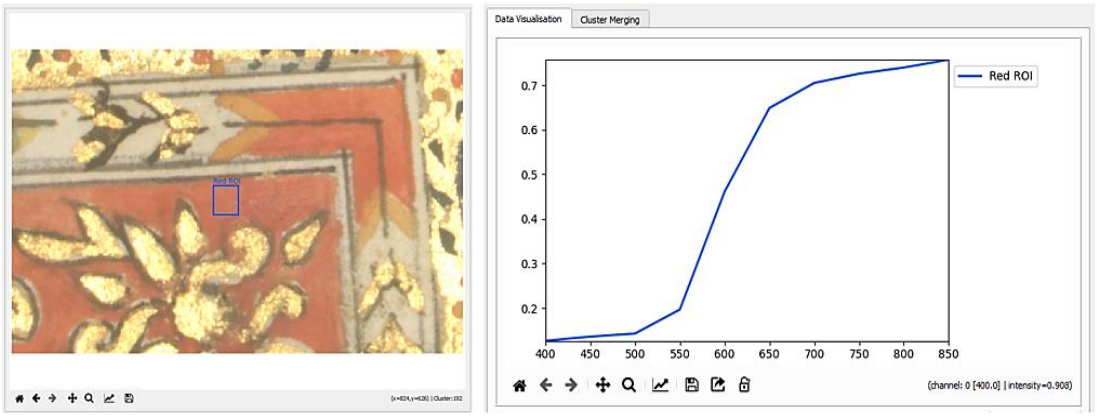

Fig.6.4 A representation of performing spectral extraction in Guisi. The left-hand side shows an RGB image with an ROI selected, and the right-hand side shows a graph of the mean reflectance spectrum in the ROI area.

This process of spectral extraction and visualisation is one of the most essential parts of Guisi for the analysis of spectral imaging data in any study. This is because the data plotted directly from clusters or ROIs and shown in the spectral extraction window can be exported from the software into a *.csv file which subsequently can be used anywhere. In the context of this thesis for example, *.csv files for vis-NIR PRISMS data could be extracted from the software and used as unknown spectra within the Kubelka-Munk model for the identification of pigments, as is demonstrated in abundance in chapter 3.

6.3.1.3 Cluster Visualisation

Cluster visualisation allows for cluster maps to be visualised in the bottom right-hand window of Guisi, a copy of which will be shown in the main display if the user chooses. The different options for visualising certain variations of the cluster map (shown in Fig. 6.5) are as follows:

- Full Greyscale Cluster Map – a representation of all the clusters can be displayed in greyscale, where the pixel intensity corresponds to the cluster number.
- Binary Cluster Map - A black-and-white binary cluster map can be obtained for specific cluster numbers which can be selected by a user.
- RGB Cluster Map – The RGB cluster map is essentially the same as the binary cluster map but also allows for the true colour representation of the cluster pixels by using the binary map as a mask.

Fig.6.5 Different images which can be found displayed within the cluster visualisation window of Guisi. From Left: sRGB Image, Greyscale cluster map, binary cluster map for vermillion, true colour cluster map for vermillion.

Different cluster map variations can be useful for analysis in different ways. The full greyscale maps, which can also be converted to false colour, can be useful for investigating the dispersion of different clusters over a painting, with the binary and RGB cluster maps useful for mapping the distribution of pixels belonging to specific pigment mixtures either in a single painting or over an entire collection. In addition, all the cluster maps are used in cluster grouping, where smaller ROIs of cluster map regions can be used to highlight spatial areas with similar clusters (as shown in section 6.3.1.4). The visualisation of these maps was useful throughout the thesis in all the different case studies, with the RGB cluster maps seen numerous times in different chapters. Furthermore, even the classification maps produced in chapter 5 could be used within Guisi and were shown to be used to map classified pigment mixtures over the collection when evaluating the C-SVM performance. This flexibility in cluster visualisation is provided through the use of cluster maps where pixel intensities correspond to the cluster to which the pixel belongs. This implies that if the same format is followed for any clustering or analysis technique involving any form of cluster map in any data type, then the results can be used within Guisi, making it a very universal approach to data analysis – precisely is needed for holistic multimodal studies.

6.3.1.4 Cluster ROI Tools

The upper part of the cluster merging tab is used to primarily interact with spectral imaging data to gather clusters from selected spatial areas in the main display. This is typically performed either by selecting a live ROI, which can be moved around to any point on an image or by using a saved ROI created during spectral extraction. For whichever type of ROI is used, the left-hand side of the cluster ROI tools window will, as a result, show either the RGB, mean, or cluster map image of the selected region, easily switchable by the user. Within this region will exist a set of unique clusters, each of which will possess a unique mean cluster spectrum which will be displayed on the right-hand side of the window.

A typical example of this window in use for a red region of interest can be seen in Fig. 6.6, where the leftmost part of the figure. displays the greyscale cluster map, the middle the RGB for the selected region, and the rightmost part shows the mean cluster spectra belonging to the clusters within the ROI.

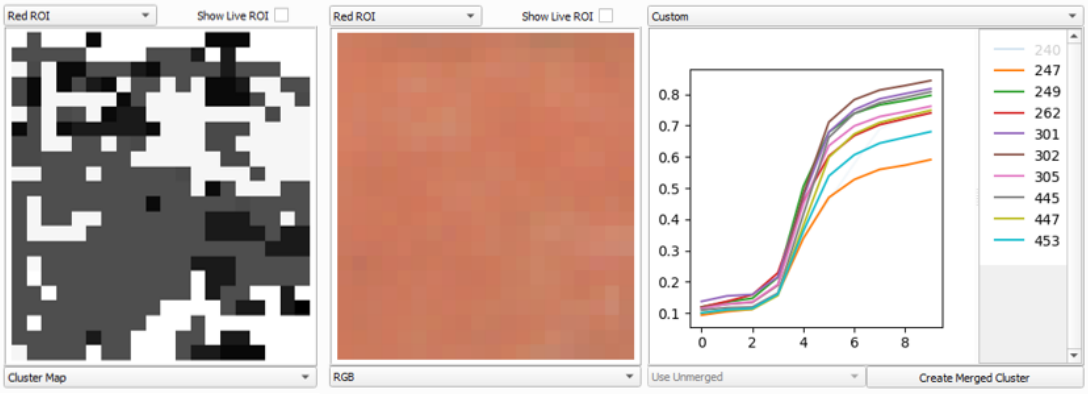

Fig.6.6 Different Guisi displays found within the cluster ROI tools window. From Left: ROI Greyscale cluster map; ROI sRGB Image; a plot of vis-NIR mean cluster spectra belonging to each unique cluster within the ROI.

After selecting a region of interest, clusters can be merged by simply clicking the "create merged cluster" button, which will assign the different clusters into a single new group, which will be sent through to the cluster grouping part of Guisi (shown in 6.3.1.5) to be further analysed, edited, and eventually used to create new group clusters maps. If a selected region of interest possesses spectra that may lay within the selected region of interest but do not possess the same features as most of the other clusters in the region, spectra can be removed from the set to be merged by clicking either on their plot or legend entry, as shown with the greyed-out cluster 240 in Fig. 6.6, ensuring accurate grouping of clusters which can clearly be shown to have the same spectrum.

6.3.1.5 Cluster Grouping

Cluster grouping can be performed by merging the clusters selected in the cluster ROI tools, or by creating an empty group and populating it with clusters afterwards. Either way, clusters can be added or removed to and from groups by using the cluster list widget and the popup window which will hold group properties (as seen in Fig. 6.7). Within each popup window, the constituent clusters belonging to each group can be edited, and as clusters are added or removed from the group, the associated cluster means are updated. Multiple options can be edited within the popup cluster options window, these include:

- Changing the cluster number: This will in turn update group cluster maps if produced. The cluster cannot be changed to a number which already exists for another cluster to avoid confusion.

- Changing the colour of the cluster: This will update not only the plot colours inside the software but also affects the colour of the false colour group cluster maps.
- Adding Clusters: Clusters can be manually added to the group if they possess similar features. Other groups can also be merged into the group being edited as well, which allows for the easier combining of cluster groups representing the same materials over multiple ROIs.
- Removing Clusters: Constituent clusters can be removed from the cluster if any anomalies are found.
- Changing Metadata: By clicking on the metadata tab, an open text area will be found where anything typed will be stored in the group for future reference. This can be useful when labelling different clusters with what they represent over large collections.

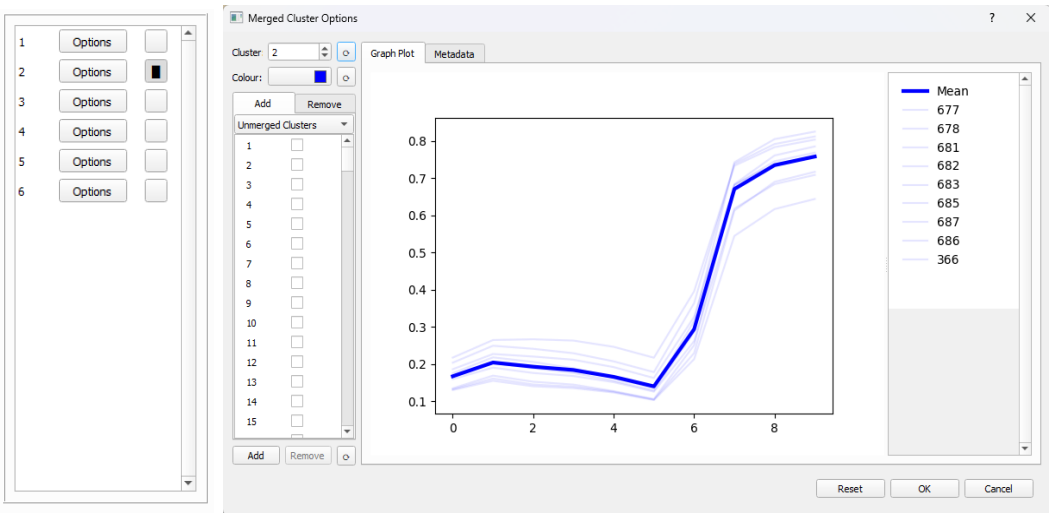

Fig.6.7 Illustration of the Guisi grouping options after opening from the panel on the left-hand side.

After analysis, grouping data can be exported which will create a new set of cluster maps and results data which will have negative cluster numbers for merged groups. This means that ungrouped clusters can be placed into the same maps as grouped clusters, though if requested just the grouped data can be produced also (which is typically how the grouping was used for the research within this study). Output cluster maps and the corresponding data can also be placed into Guisi again later if another stage of grouping is preferred. New results data is created from the clustering outputs described in chapter 2, which includes the total sum of reflectance ($\sum$R) $\sum_{i=1}^{N} R_{i,\lambda}$, the total sum of squared reflectance ($\sum$R^2) $\sum_{i=1}^{N} R_{i,\lambda}^2$, and total pixel population N, so pixel-accurate means and standard deviations representing each cluster can be calculated quickly and used for analysis or further grouping.

6.4 Conclusions & Further Developments

In this chapter, a new software application, Guisi, was introduced for the purposes of performing spectral imaging data analysis and grouping of clusters which are generated after performing the automated clustering procedures seen in chapter 2. An acquisition to results workflow (ARW) for spectral imaging studies was briefly mentioned which illustrated where Guisi could be implemented for numerous purposes, both within the studies in this thesis and in future research projects.

Questions were asked about the data formats which should be used in cultural heritage studies of large data collections, where it was discussed that many techniques can be useful in their applications, however, to perform large-scale or holistic multimodal studies with different spectral imaging data types, the commonly supported and open access BigTIFF file format is an excellent candidate for studies in this thesis and beyond. Finally, the software itself was discussed and demonstrated, where a more detailed use of the application for the extraction of spectra and grouping of clusters was shown.

Moving forward, intentional future developments of Guisi are to integrate more tools useful for the analysis of spectral data and identification of materials in cultural heritage studies. There are already other codes and basic GUI programs that have been developed separately for analytical processes such as XRF element identification or Kubelka-Munk mixture modelling for the identification of pigments, respectively. However, as of yet, these processes are developed as separate algorithms or exist in other programs such as MATLAB. By bringing these tools into the software, Guisi can begin to become a platform in which large-scale complex data analysis can be performed accurately, efficiently, and conveniently. Furthermore, with the results shown in various chapters within this thesis, it is possible that in the future, clustering and classification techniques could be implemented directly into the software also, so that in future Guisi can be used to analyse multimodal spectral imaging data and large collections throughout all steps of the proposed ARW.

Finally, Guisi is being partly developed as part of the “AI for DIGILAB” (https://www.isaac-lab.com/isaac-digilab) project and is likely to adapt, improve and develop as researchers from different cultural institutions or other organisations begin to use the available services involved with the project. Therefore, it is also the intention that in the future multiple other tools useful for the interpretation of spectral data not already mentioned within this thesis will be incorporated into the software, so that Guisi can grow and then continue to be a dynamic software application useful throughout the broader cultural heritage research community.

Chapter 7
Conclusions

The main objective of this thesis was to develop and implement working methodologies which would allow for the large-scale analysis of complex spectral imaging data so that the British Library's collection of Southeast Asian illuminated manuscripts could undergo efficient, thorough and widespread material analysis. The main motivations for this project were to acquire a fundamental basis of knowledge surrounding the evolution of trade and cultural exchanges within maritime Southeast Asia by studying how the development of artistic practises and use of artistic materials in the production of illuminated manuscripts may have adapted over time and between regions. While maritime Southeast Asia has a well-documented history of producing finely illuminated manuscripts, very little is known about the artistic practises implemented within their production throughout the 18^{th}-19^{th} centuries. Furthermore, as the number of well-illuminated manuscripts is limited, and the use of artistic materials largely unknown, basic comparative studies may not be informative or statistically reliable enough to provide conclusions about the materials used in the creation of illuminated manuscripts. Therefore, the holistic use of reflectance spectral imaging with other complementary techniques was considered a perfect candidate methodology for performing a large-scale material analysis. However, when performing studies covering vast geographical regions, a large volume of data is required to make confident conclusions about the use of different artistic materials, resulting in data which can become too large or complex to analyse. Therefore, it was determined that machine learning techniques were required to confidently characterise artistic materials with significant variation spread over a large amount of collection items.

Prior to the beginning of this research project, the large-scale analysis of spectral imaging data had not been explored thoroughly within cultural heritage science. However, one clustering approach had already been explored and implemented by our research group for the large-scale material identification of wall paintings (Kogou et al., 2020). While this technique had performed well for the application outlined in that study, expanding the method to alternative datasets showed that the algorithm suffered from multiple difficulties where data would not be clustered correctly. Therefore, the initial focus of the research outlined in this thesis was to develop a generalised machine-learning methodology which would allow for accurate material identification to be performed on the British Library's large and spectrally complex illuminated manuscript collection. The developed method eventually came to be a novel fully automated complementary multimodal clustering technique (multimodal meaning more than one type of spectroscopy/data type corresponding to the same object) which could perform large-scale clustering of collections by using VNIR spectral imaging data and CIELAB colour information

together in a unified way, which demonstrated the ability to successfully cluster previously misclustered data. For performing analysis on the British Library's maritime Southeast Asian illuminated manuscript collection, the automated clustering was able to provide a reduction in the size of the dataset by roughly a factor of 84,200, meaning that a collection of almost 2 billion spectra could be reduced to 22,837 unique clusters representing different artistic materials including variations in substrate, inks, pigments, pigment mixtures, paint layers, and more.

While this methodology was shown to provide accurate clustering results which could demonstrate the presence of the same materials spread over entire collections, analysing 22,837 individual clusters for this thesis was not an ideal approach for performing large-scale data analysis. Therefore, a highly reliable method of placing clusters into pigment mixture groups was also developed and implemented which used the Kubelka-Munk model for the identification of pigments to inform the grouping of clusters which shared the same pigment mixtures, regardless of relative concentration between individual constituent materials. As this process would be highly accurate in theory but difficult to implement in practise, a new software application, named Guisi (a Graphical User Interface for Spectral Imaging and Clustering Analysis), was developed so that the spectral data could be analysed more efficiently and automated clustering results could be assigned into groups by interacting directly VNIR data, colour data, and clustering results simultaneously through convenient data visualisation tools. In the context of this thesis, this meant that grouping could be used to illustrate the distribution of pigments and pigment mixtures for over a thousand spectral images by interacting with only a small set of data cubes, making accurate large-scale material analysis more achievable for many different items within the British Library's collection.

When analysing these automated clustering results and grouping results over the entire manuscript collection during pigment identification for the whole collection, it became obvious that there were fundamental limitations which led to the misclustering of spectrally similar pigments and materials which in truth possessed chemical differences. These limitations were mainly a result of performing clustering in only the VNIR, however, for the analysis of the collection this was not a huge problem, as it was found that by implementing other complementary point-based measurement techniques, such as XRF and ER-FTIR, many of the misclustered regions could be easily identified and then correctly characterised anyway. This did however make it clear that the natural progression of performing automated clustering for holistic studies was to include more than just one technique in the clustering approach itself. This is because the complementary information available via imaging variants of the point analysis techniques can provide data to not only identify specific material composition, but also map different materials and mixtures more accurately.

Therefore, the automated clustering method was investigated once again to determine its effectiveness in clustering alternative complementary spectral imaging techniques, where it was found that the same clustering algorithm developed for VNIR could also be used to cluster XRF and ER-FTIR data. Furthermore, as the automated clustering method developed for CIELAB/VNIR spectral imaging could allow for two datasets to be clustered at a time, an adaptation of the algorithm was created to allow for the automated clustering of all techniques simultaneously, where for two export paintings, XRF mapping and ER-FTIR mapping were used directly with VNIR spectroscopy and CIELAB within the clustering. This novel holistic multimodal clustering approach showed success in addressing the misclustering problems seen commonly in each individual technique, as the inclusion of complementary information improved the overall accuracy. In addition to this however, the approach also provided cluster maps which could provide a holistic analysis of artistic materials and demonstrate the layering of different paints by combining both the XRF and VNIR. Furthermore, in general, the clustering approach was also able to provide much needed higher SNR spectra required for XRF (where it is often preferable to collect data quickly instead of using longer spectra collection times for higher SNR) so that more accurate pigment identification could be performed.

After performing automated clustering on other projects outside this thesis' scope, along with the dataset used during the investigation into holistic multimodal clustering, it was discovered that the same pigment mixtures used within the British Library's maritime Southeast Asian illuminated manuscript collection could often appear in other analysed artworks. Multiple successes in performing automated clustering on other datasets presented the question of whether results gained from one study could be used to inform the study of another completely different small or large collection, in turn bypassing much of the initial grouping that may be required with brand new datasets. As machine learning had already been implemented in the automated clustering and grouping methodology used to analyse large collections, it was investigated again to see whether classification techniques could be used to achieve this new objective. In response, several different commonly implemented classification techniques were investigated in this thesis to understand if they could be used to automatically classify previously encountered and characterised VNIR spectra. After training different models on a small subset of manuscripts from the British Library's collection and applying them to the rest of the dataset, it was found that classification can, for the most part, successfully classify already seen pigment mixtures in new collections. This prompted the development of a new methodology which will be investigated in the future, which involves a unified clustering, grouping, and classification workflow that can allow for automated pigment identification to be performed for already seen pigment mixtures, while also allowing misclassifications to be used as further inputs into the KM informed grouping, thus producing a loop which will iteratively update and improve as more materials are identified, classified and trained into the classification model.

Finally, with regards to the pigment identification performed for the British Library's collection of maritime Southeast Asian illuminated manuscripts, the different clustering and grouping techniques were shown to perform well in mapping pigments over the entire collection when using the KM informed approach. These results, used in combination with UV-VNIR-SWIR reflectance, XRF spectroscopy, ER-FTIR spectroscopy and Raman spectroscopy were able to perform identification of many pigments and mixtures spread wide throughout numerous manuscripts with provenances ranging over two centuries. While the findings are explained in detail in chapter 3, some of the overarching trends which could be determined are as follows:

- As very little was known about any pigment usage, the identification of lead white, bone white, calcium carbonate white, barium white, indigo, Prussian blue, ultramarine, azurite (potentially), vermillion, red lead (potentially), red insect-based dyes, red ochre, orpiment, yellow ochre, chrome yellow, gamboge, copper greens, iron gall ink, copper-based ink, carbon black (likely) and bone black (potentially) were all detected within the collection.
- Most non-primary colours such as greens, oranges, browns, pinks etc. were likely to almost always be produced in manuscripts by mixing the same materials used elsewhere on the manuscripts, unless copper-based green or brown ochre was used.
- The previously unmentioned use of bone ash white in early 19th century manuscripts was very common in Java and can likely be used to date some other bone ash containing manuscripts (Add MS 12287, Add MS 12292, Add MSS 12339) with unspecified dates.
- In general, it was found that until the mid-late 19th century, most regions within maritime Southeast Asia used relatively similar materials within their manuscript illumination that they had used since at least the mid-18th century. The changes in the mid-late 19th century however were seen most clearly for blue, yellow and white pigments, where towards the later dates, the use of more traditional materials such as indigo, orpiment, and bone ash white had changed to Prussian blue, ultramarine, chrome yellow of barium white: see Or 15227 (Late 19th century), Or 15026 (dated to 1861) and Or 9333 (dated late 19th-early 20th century).
- During these two centuries, it seems the greatest "new" influences were likely European, where materials such as Prussian Blue, Ultramarine (likely synthetic), Chrome Yellow, and Barium white were introduced in most in manuscripts which typically had less traditional features.

In addition to the future developments already mentioned regarding the unified clustering, classification, and pigment identification workflow, there are also other future developments which will be explored due to the positive results seen in this thesis.

Firstly, as in general most of the clustering throughout this thesis is performed using the same automated method developed in chapter 2, it would be interesting in future to continue investigating how well this generalised automated clustering algorithm would work with alternative spectral imaging systems that may use different spectral ranges or have greater focus on different parts of the VNIR, SWIR etc. It would also be interesting to investigate new distance metrics at the clustering stage to understand if the approach can be optimised for different spectral resolutions or techniques different from PRISMS.

Chapter 3 demonstrated clearly that the clustering techniques and overall methodology could perform well for the investigation of the maritime Southeast Asian illuminated manuscript collection. While there is significant complexity to the variation and application of different pigments found in the collection, there are perhaps some general scenarios in different studies where an even greater level of complexity could be encountered. While the clustering method has already been implemented on different sets of data including VNIR spectral imaging/CIELAB of Columbian export paintings, and XRF/SWIR hyperspectral imaging of old master paintings, it would be interesting to continue investigating how wide the range of analysis using this technique can be, and if this generalised clustering algorithm can in fact work for all spectral imaging-based heritage science investigations.

With the successes of holistic multimodal clustering using VNIR spectral imaging/CIELAB/XRF together were clear within this thesis, it would be interesting in future to investigate whether the holistic multimodal clustering of collections can be performed. Or if even more, higher-spatial resolution techniques can be incorporated into the clustering approach. This is because if the same approach can be extended to many high-resolution non-invasive imaging and mapping techniques, such as VNIR, XRF, and FTIR, a completely automated holistic analysis of artworks could be performed without losing information during image registration between different techniques.

Finally, while Guisi demonstrated its importance for the different analysis and grouping implementations seen in different chapters, there are still many other data analysis and processing techniques, such as Kubelka-Munk modelling for pigment identification, which are not integrated into the software application at the time of writing this thesis. Therefore, there are obvious future developments which can still be made to the software so that Guisi can one day become an open platform for multiple spectral imaging techniques in cultural heritage.

Bibliography

Aceto, M., Agostino, A., Fenoglio, G., Idone, A., Crivello, F., Griesser, M., Kirchweger, F., Uhlir, K., & Puyo, P. R. (2017). Analytical investigations on the Coronation Gospels manuscript. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, *171*, 213–221. https://doi.org/10.1016/J.SAA.2016.07.050

Adelabu, S., Mutanga, O., & Adam, E. (2015). Testing the reliability and stability of the internal accuracy assessment of random forest for classifying tree defoliation levels using different validation methods. *Geocarto International*, *30*(7), 810–821. https://doi.org/10.1080/10106049.2014.997303

Aga Khan Museum. (2023). *Qur'an Manuscript, AKM488, The Aga Khan Museum*. https://www.agakhanmuseum.org/collection/artifact/quran-akm488

AGTA. (2023). *Lapis Lazuli | AGTA*. https://agta.org/education/gemstones/lapis-lazuli/

Alberti, R., Frizzi, T., Bombelli, L., Gironda, M., Aresi, N., Rosi, F., Miliani, C., Tranquilli, G., Talarico, F., & Cartechini, L. (2017). CRONO: a fast and reconfigurable macro X-ray fluorescence scanner for in-situ investigations of polychrome surfaces. *X-Ray Spectrometry*, *46*(5), 297–302. https://doi.org/https://doi.org/10.1002/xrs.2741

Alfeld, M., Janssens, K., Dik, J., De Nolf, W., & Van Der Snickt, G. (2011). Optimization of mobile scanning macro-XRF systems for the in situ investigation of historical paintings. *Journal of Analytical Atomic Spectrometry*, *26*(5), 899–909. https://doi.org/10.1039/c0ja00257g

Alfeld, M., Pedetti, S., Martinez, P., & Walter, P. (2018). Joint data treatment for Vis–NIR reflectance imaging spectroscopy and XRF imaging acquired in the Theban Necropolis in Egypt by data fusion and t-SNE. *Comptes Rendus Physique*, *19*(7), 625–635. https://doi.org/https://doi.org/10.1016/j.crhy.2018.08.004

Arjasakusuma, S., Swahyu Kusuma, S., & Phinn, S. (2020). Evaluating Variable Selection and Machine Learning Algorithms for Estimating Forest Heights by Combining Lidar and Hyperspectral Data. *ISPRS International Journal of Geo-Information*, *9*(9). https://doi.org/10.3390/ijgi9090507

Aru, M., Burgio, L., & Rumsey, M. S. (2014). Mineral impurities in azurite pigments: Artistic or natural selection? *Journal of Raman Spectroscopy*, *45*(11–12), 1013–1018. https://doi.org/10.1002/jrs.4469

Babić, R. J. (2022). *Cultural heritage image classification using transfer learning for feature extraction: a comparison*.

Balaram, P. T. (2012). *"Indian Indigo" in The Materiality of Color: The production, circulation, and application of dyes and pigments, 1400-1800.*

Barnett, J. R., Miller, S., & Pearce, E. (2006). Colour and art: A brief history of pigments. *Optics & Laser Technology*, *38*(4), 445–453. https://doi.org/https://doi.org/10.1016/j.optlastec.2005.06.005

Bartoll, J. (2008). The early use of Prussian blue in paintings. *Proceedings of the 9th International Conference on NDT of Art*.

Biron, C., Mounier, A., Arantegui, J. P., Bourdon, G. Le, Servant, L., Chapoulie, R., Roldán, C., Almazán, D., Díez-de-Pinos, N., & Daniel, F. (2020). Colours of the « images of the floating world ». non-invasive analyses of Japanese ukiyo-e woodblock prints (18th and 19th centuries) and new contributions to the insight of oriental materials. *Microchemical Journal*, *152*, 104374. https://doi.org/https://doi.org/10.1016/j.microc.2019.104374

Bruker. (n.d.). *ESPRIT Spectrum | Bruker*. Retrieved March 20, 2023, from https://www.bruker.com/en/products-and-solutions/elemental-analyzers/eds-wds-ebsd-SEM-Micro-XRF/software-esprit-family/esprit-spectrum.html

Bruker. (2023). *M6 JETSTREAM | Bruker*. https://www.bruker.com/en/products-and-solutions/elemental-analyzers/micro-xrf-spectrometers/m6-jetstream.html

Burgio, L., Clark, R. J. H., & Gibbs, P. J. (1999). Pigment identification studies in situ of Javanese, Thai, Korean, Chinese and Uighur manuscripts by Raman microscopy. *Journal of Raman Spectroscopy*, *30*(3), 181–184. https://doi.org/10.1002/(SICI)1097-4555(199903)30:3<181::AID-JRS356>3.0.CO;2-8

Butler, L. R. P. (2018). *Modeling cumulative effects of heavy metal contamination in mining areas of the Rimac basin*.

Caggiani, M. C., & Colomban, P. (2018). *Raman microspectroscopy for Cultural Heritage studies*. *3*(11). https://doi.org/doi:10.1515/psr-2018-0007

Caliri, C., Bicchieri, M., Biocca, P., & Romano, F. P. (2021). In situ macro X-Ray fluorescence scanning on a Leonardo da Vinci portrait. *X-Ray Spectrometry*, *50*(4), 332–340. https://doi.org/10.1002/xrs.3193

CAMEO. (2020, November 3). *Bone White*. https://cameo.mfa.org/wiki/Bone_white

Casadio, F., Daher, C., & Bellot-Gurlet, L. (2016). Raman Spectroscopy of cultural heritage Materials: Overview of Applications and New Frontiers in Instrumentation, Sampling Modalities, and Data Processing. *Topics in Current Chemistry*, *374*(5), 62. https://doi.org/10.1007/s41061-016-0061-z

Chakravarty, S., Paikaray, B. K., Mishra, R., & Dash, S. (2021). Hyperspectral Image Classification using Spectral Angle Mapper. *2021 IEEE International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE)*, 87–90. https://doi.org/10.1109/WIECON-ECE54711.2021.9829585

Chang, N. Bin, & Bai, K. (2018). Multisensor data fusion and machine learning for environmental remote sensing. *Multisensor Data Fusion and Machine Learning for Environmental Remote Sensing*, *January*, 1–508. https://doi.org/10.1201/b20703

Chaplot, S., Patnaik, L. M., & Jagannathan, N. R. (2006). Classification of magnetic resonance brain images using wavelets as input to support vector machine and neural network. *Biomedical Signal Processing and Control*, *1*(1), 86–92. https://doi.org/https://doi.org/10.1016/j.bspc.2006.05.002

Cie. (2004). Cie 15: Technical Report. *Colorimetry, 3rd Edition*, *552*, 24. https://doi.org/ISBN 3 901 906 33 9

Clark, R. J. H., Cooksey, C. J., Daniels, M. A. M., & Withnall, R. (1993). Indigo, woad, and Tyrian Purple: important vat dyes from antiquity to the present. *Endeavour*, *17*(4), 191–199. https://doi.org/https://doi.org/10.1016/0160-9327(93)90062-8

Coccato, A., Moens, L., & Vandenabeele, P. (2017). On the stability of mediaeval inorganic pigments: a literature review of the effect of climate, material selection, biological activity, analysis and conservation treatments. *Heritage Science*, *5*(1), 12. https://doi.org/10.1186/s40494-017-0125-6

Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273–297. https://doi.org/10.1007/BF00994018

Cottrell, M., Olteanu, M., Rossi, F., Villa-Vialaneix, N. N., & Villa-Vialaneix, N. (2018). Self-Organizing Maps, theory and applications. In *Revista de Investigacion Operacional* (Vol. 39, Issue 1). https://hal.science/hal-01796059

Cover, T., & Hart, P. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, *13*(1), 21–27. https://doi.org/10.1109/TIT.1967.1053964

Cucci, C., Delaney, J. K., & Picollo, M. (2016). Reflectance Hyperspectral Imaging for Investigation of Works of Art: Old Master Paintings and Illuminated Manuscripts. *Accounts of Chemical Research*, *49*(10), 2070–2079. https://doi.org/10.1021/acs.accounts.6b00048

Cupitt, J., & Martinez, K. (1996). VIPS: An imaging processing system for large images. *Proceedings of SPIE - The International Society for Optical Engineering*, *1663*. https://doi.org/10.1117/12.233043

Delaney, J. K., Ricciardi, P., Glinsman, L. D., Facini, M., Thoury, M., Palmer, M., & Rie, E. R. de la. (2014). Use of imaging spectroscopy, fiber optic reflectance spectroscopy, and X-ray fluorescence to map and identify pigments in illuminated manuscripts. *Studies in Conservation*, *59*(2), 91–101. https://doi.org/10.1179/2047058412Y.0000000078

Dooley, K. A., Conover, D. M., Glinsman, L. D., & Delaney, J. K. (2014). Complementary Standoff Chemical Imaging to Map and Identify Artist Materials in an Early Italian Renaissance Panel Painting. *Angewandte Chemie International Edition*, *53*(50), 13775–13779. https://doi.org/https://doi.org/10.1002/anie.201407893

Duffy, C. (2018). Multi-spectral Imaging at the British Library. *2018 3rd Digital Heritage International Congress (DigitalHERITAGE) Held Jointly with 2018 24th International Conference on Virtual Systems & Multimedia (VSMM 2018)*, 1–4. https://doi.org/10.1109/DigitalHeritage.2018.8810072

Eastaugh, N., Walsh, V., Chaplin, T., & Siddall, R. (2008). *The Pigment Compendium A dictionary of historical pigments (Bone, Calcined) pg. 58* (1st Edition). Routledge.

Eremin, K., Stenger, J., Huang, J.-F., Aspuru-Guzik, A., Betley, T., Vogt, L., Kassal, I., Speakman, S., & Khandekar, N. (2008). Examination of pigments on Thai manuscripts: the first identification of copper citrate. *Journal of Raman Spectroscopy*, *39*(8), 1057–1065. https://doi.org/https://doi.org/10.1002/jrs.1985

Fan, C., Zhang, P., Wang, S., & Hu, B. (2018). A study on classification of mineral pigments based on spectral angle mapper and decision tree. *Proc.SPIE*, *10806*, 108065Z. https://doi.org/10.1117/12.2503088

Feller, R. L. (1986). *Artists pigments : a handbook of their history and characteristics. vol. 1.* (Vol. 1). Cambridge University Press.

Fichera, G. V., Malagodi, M., Cofrancesco, P., Weththimuni, M. L., Guglieri, C., Olivi, L., Ruffolo, S., & Licchelli, M. (2018). Study of the copper effect in iron-gall inks after artificial ageing. *Chemical Papers*, *72*(8), 1905–1915. https://doi.org/10.1007/s11696-018-0412-z

Fiorucci, M., Khoroshiltseva, M., Pontil, M., Traviglia, A., Del Bue, A., & James, S. (2020). Machine Learning for Cultural Heritage: A Survey. *Pattern Recognition Letters*, *133*, 102–108. https://doi.org/https://doi.org/10.1016/j.patrec.2020.02.017

Fischer B, Smith M, & Pau G. (2022). *rhdf5: R Interface to HDF5* (version 2.42.0). R package. https://github.com/grimbough/rhdf5.

FitzHugh, E. W., Leona, M., Winter, J., & Freer Gallery of Art. (2003). Pigments in later Japanese paintings New series Vol. 1. In *2003*. https://library.si.edu/digital-library/book/pigmentsinlaterj12003fitz

Fix, E., & Hodges, J. L. (1951). *Discriminatory analysis, nonparametric discrimination*.

Foody, G. M., & Mathur, A. (2004). A relative evaluation of multiclass image classification by support vector machines. *IEEE Transactions on Geoscience and Remote Sensing*, *42*(6), 1335–1343. https://doi.org/10.1109/TGRS.2004.827257

Forgy, E. W. (1965). Cluster analysis of multivariate data : efficiency versus interpretability of classifications. *Biometrics*, *21*, 768–769.

French, K., & Monaghan, M. (2018). PRESERVING THE SOUTHEAST ASIAN PAINTINGS AT THE WALTERS ART MUSEUM. *The Journal of the Walters Art Museum*, *73*, 108–112. https://www.jstor.org/stable/26537588

Fukushima, K. (1975). Cognitron: A self-organizing multilayered neural network. *Biological Cybernetics*, *20*(3–4), 121–136. https://doi.org/10.1007/BF00342633

Galeotti, M., Mazzeo, R., & Pinna, D. (viaf)162553086. (2009). *Scientific examination for the investigation of paintings : a handbook for conservators-restorers*. Florens : Centro Di. http://lib.ugent.be/catalog/rug01:001401768

Gallop, A. (2004). An Acehnese Style of Manuscript Illumination. *Archipel*, 193–240. https://www.persee.fr/doc/arch_0044-8613_2004_num_68_1_3834

Gallop, A. (2013). *The Crown of Kings: a deluxe Malay manuscript from Penang - Asian and African studies blog*. https://blogs.bl.uk/asian-and-african/2013/09/the-crown-of-kings-a-deluxe-malay-manuscript-from-penang.html

Gallop, A. (2015). *Royal genealogies from Indonesia and the Malay world - Asian and African studies blog*. https://blogs.bl.uk/asian-and-african/2015/04/royal-genealogies-from-indonesia-and-the-malay-world-.html

Gallop, A. (2022). *Frederik Jacob Rothenbühler and his wife as collectors of Javanese manuscripts in the early 19th century - Asian and African studies blog*. https://blogs.bl.uk/asian-and-african/2022/09/frederik-jacob-rothenbuhler-and-his-wife-as-collectors-of-javanese-manuscripts-in-the-early-19th-cen.html

Gamarra-Urrunaga, J. E., Castroviejo, R., & Bernhardt, H.-J. (2013). Preliminary mineralogy and ore petrology of the intermediate-sulfidation Pallancata Deposit, Ayacucho, Peru. *The Canadian Mineralogist*, *51*(1), 67–91. https://doi.org/10.3749/canmin.51.1.67

Gao, Z., Du, M., Cao, N., Hou, M., Wang, W., & Lyu, S. (2023). Application of hyperspectral imaging technology to digitally protect murals in the Qutan temple. *Heritage Science*, *11*(1), 1–16. https://doi.org/10.1186/S40494-022-00847-7/FIGURES/15

Gao, Z., Shao, Y., Xuan, G., Wang, Y., Liu, Y., & Han, X. (2020). Real-time hyperspectral imaging for the in-field estimation of strawberry ripeness with deep learning. *Artificial Intelligence in Agriculture*, *4*, 31–38. https://doi.org/10.1016/J.AIIA.2020.04.003

Ghamisi, P., Rasti, B., Yokoya, N., Wang, Q., Hofle, B., Bruzzone, L., Bovolo, F., Chi, M., Anders, K., Gloaguen, R., Atkinson, P. M., & Benediktsson, J. A. (2019). Multisource and Multitemporal Data Fusion in Remote Sensing: A Comprehensive Review of the State of the Art. *IEEE Geoscience and Remote Sensing Magazine*, *7*(1), 6–39. https://doi.org/10.1109/MGRS.2018.2890023

Goshtasby, A. (1986). Piecewise linear mapping functions for image registration. *Pattern Recognition*, *19*(6), 459–466. https://doi.org/https://doi.org/10.1016/0031-3203(86)90044-0

Goshtasby, A. (1988). Image registration by local approximation methods. *Image and Vision Computing*, *6*(4), 255–261. https://doi.org/https://doi.org/10.1016/0262-8856(88)90016-9

Grabowski, B., Masarczyk, W., Głomb, P., & Mendys, A. (2018). Automatic pigment identification from hyperspectral data. *Journal of Cultural Heritage*, *31*, 1–12. https://doi.org/10.1016/J.CULHER.2018.01.003

Gumansalangi, F., Calle, J. L. P., Barea-Sepúlveda, M., Manikharda, Palma, M., Lideman, Rafi, M., Ningrum, A., & Setyaningsih, W. (2023). A Rapid Method for Authentication of Macroalgae Based on Vis-NIR Spectroscopy Data Combined with Chemometrics Approach. *Water (Switzerland)*, *15*(1). https://doi.org/10.3390/w15010100

h5py. (n.d.). *HDF5 for Python — h5py 3.8.0 documentation*. Retrieved March 20, 2023, from https://docs.h5py.org/en/stable/

Hamerton, I., Tedaldi, L., & Eastaugh, N. (2013). A Systematic Examination of Colour Development in Synthetic Ultramarine According to Historical Methods. *PLOS ONE*, *8*(2), e50364-. https://doi.org/10.1371/journal.pone.0050364

Han, B., Zhang, B., Chong, J., Sun, Z., & Yang, Y. (2022). Beauty and chemistry: the independent origins of synthetic lead white in east and west Eurasia. *Humanities and Social Sciences Communications*, *9*(1), 281. https://doi.org/10.1057/s41599-022-01290-6

Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A K-Means Clustering Algorithm. In *Source: Journal of the Royal Statistical Society. Series C (Applied Statistics)* (Vol. 28, Issue 1).

Hashemi-Nasab, F. S., & Parastar, H. (2022). Vis-NIR hyperspectral imaging coupled with independent component analysis for saffron authentication. *Food Chemistry*, *393*. https://doi.org/10.1016/j.foodchem.2022.133450

Havermans, J., Aziz, H., & Scholten, H. (2003). Non Destructive Detection of Iron Gall Inks by Means of Multispectral Imaging Part 1: Development of the Detection System. *Restaurator-International Journal for The Preservation of Library and Archival Material - RESTAURATOR, 24*, 55–60. https://doi.org/10.1515/REST.2003.55

IEC. (1999). *Multimedia systems and equipment - Colour measurement and management - Part 2-1: Colour management - Default RGB colour space - sRGB.*

Jensen, E. (2023). *Ultramarine Blue Pigment, Synthetic*. https://www.mccrone.com/mm/ultramarine-blue-pigment-synthetic/

Kantoğlu, Ö., Ergun, E., Kırmaz, R., Kalaycı, Y., Zararsız, A., & Bayır, Ö. (2018). *Colour and Ink Characterization of Ottoman Diplomatic Documents Dating from the 13th to the 20th Century*. *39*(4), 265–288. https://doi.org/doi:10.1515/res-2018-0014

Kim, N. (2020). *Orpiment 石黃, also 雄黃 in Yunnan*. https://www.zo.uni-heidelberg.de/sinologie/research/mining-sw/minerals/orpiment.html

Kleynhans, T., Patterson, C., Dooley, K., Messinger, D., & Delaney, J. (2020). *An Alternative Approach to Mapping Pigments in Paintings With Hyperspectral Reflectance Image Cubes Using Artificial Intelligence*. https://doi.org/10.21203/rs.3.rs-40103/v1

Knipe, P., Eremin, K., Walton, M., Babini, A., & Rayner, G. (2018). Materials and techniques of Islamic manuscripts. *Heritage Science*, *6*(1), 55. https://doi.org/10.1186/s40494-018-0217-y

Kogou, S. (2017). Investiation of the complementary use of non-invasive techniques for the holistic analysis of paintings and automatic analysis of large scale spectral imaging data, https://irep.ntu.ac.uk/id/eprint/32752/

Kogou, S., Lee, L., Shahtahmassebi, G., & Liang, H. (2021). A new approach to the interpretation of XRF spectral imaging data using neural networks. *X-Ray Spectrometry*, *50*(4), 310–319. https://doi.org/https://doi.org/10.1002/xrs.3188

Kogou, S., Lucian, A., Bellesia, S., Burgio, L., Bailey, K., Brooks, C., & Liang, H. (2015). A holistic multimodal approach to the non-invasive analysis of watercolour paintings. *Applied Physics A, 121*(3), 999–1014. https://doi.org/10.1007/s00339-015-9425-4

Kogou, S., Neate, S., Coveney, C., Miles, A., Boocock, D., Burgio, L., Cheung, C. S., & Liang, H. (2016). The origins of the Selden map of China: Scientific analysis of the painting

materials and techniques using a holistic approach. *Heritage Science*, *4*(1), 1–24. https://doi.org/10.1186/s40494-016-0098-x

Kogou, S., Shahtahmassebi, G., Lucian, A., Liang, H., Shui, B., Zhang, W., Su, B., & van Schaik, S. (2020). From remote sensing and machine learning to the history of the Silk Road: large scale material identification on wall paintings. *Scientific Reports*, *10*(1), 19312. https://doi.org/10.1038/s41598-020-76457-9

Kohonen, T. (1990). The Self-organizing Map. *Proceedings of the IEEE*, *78*(9), 1464–1480. https://doi.org/10.1109/5.58325

Kohonen, T. (2001). *Self Organizing Maps* (3rd ed.). Springer-Verlag Berlin Heidelberg.

Kopplin, M. (2002). Lacquerware in Asia: China, Korea, Japan and the Ryukyu Islands. In *UNSECO Publishing*. UNESCO. https://unesdoc.unesco.org/notice?id=p::usmarcdef_0000125944

Kruse, F. A., Lefkoff, A. B., Boardman, J. W., Heidebrecht, K. B., Shapiro, A. T., Barloon, P. J., & Goetz, A. F. H. (1993). The spectral image processing system (SIPS)—interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment*, *44*(2), 145–163. https://doi.org/https://doi.org/10.1016/0034-4257(93)90013-N

Kubelka, P. (1948). New Contributions to the Optics of Intensely Light-Scattering Materials. Part I. *J. Opt. Soc. Am.*, *38*(5), 448–457. https://doi.org/10.1364/JOSA.38.000448

Kubelka, P., & Munk, F. (1931). An Article on Optics of Paint Layers (engl. Übersetzung). *Z. Tech. Phys*, *12*(1930), 593–601. http://www.graphics.cornell.edu/~westin/pubs/kubelka.pdf

Kumar, A., & McGlynn, J. (1996). *Illuminations The Writing Traditions of Indonesia* (1st ed.). Weatherhill Inc.

Kumar, R., & Sharma, V. (2017). A novel combined approach of diffuse reflectance UV–Vis-NIR spectroscopy and multivariate analysis for non-destructive examination of blue ballpoint pen inks in forensic application. *Spectrochimica Acta - Part A: Molecular and Biomolecular Spectroscopy*, *175*, 67–75. https://doi.org/10.1016/j.saa.2016.12.008

Lahat, D., Adali, T., & Jutten, C. (2015). Multimodal Data Fusion: An Overview of Methods, Challenges, and Prospects. *Proceedings of the IEEE*, *103*(9), 1449–1477. https://doi.org/10.1109/JPROC.2015.2460697

Langley, M., & O'Connor, S. (2018). *Exploring red ochre use in Timor-Leste and surrounds: Headhunting, burials, and beads*. https://doi.org/10.4324/9781315299112-3

Legrand, S., Alfeld, M., Vanmeert, F., De Nolf, W., & Janssens, K. (2014). Macroscopic Fourier transform infrared scanning in reflection mode (MA-rFTIR), a new tool for

chemical imaging of cultural heritage artefacts in the mid-infrared range. *Analyst*, *139*(10), 2489–2498. https://doi.org/10.1039/C3AN02094K

Lemay, M.-F. (2013). *Iron Gall Ink | Traveling Scriptorium*. https://travelingscriptorium.com/2013/03/21/iron-gall-ink/

Leona, M., & Winter, J. (2001). Fiber Optics Reflectance Spectroscopy: A Unique Tool for the Investigation of Japanese Paintings. *Studies in Conservation*, *46*(3), 153–162. https://doi.org/10.2307/1506806

Li, Y., Cheung, C. S., Kogou, S., Liggins, F., & Liang, H. (2019). Standoff Raman spectroscopy for architectural interiors from 3-15 m distances. *Opt. Express*, *27*(22), 31338–31347. https://doi.org/10.1364/OE.27.031338

Liang, H. (2012). Advances in multispectral and hyperspectral imaging for archaeology and art conservation. *Applied Physics A: Materials Science and Processing*, *106*(2), 309–323. https://doi.org/10.1007/s00339-011-6689-1

Liang, H., Lucian, A., Lange, R., Cheung, C., & Su, B. (2014). Remote spectral imaging with simultaneous extraction of 3D topography for historical wall paintings. *ISPRS Journal of Photogrammetry and Remote Sensing*, *95*, 13–22. https://doi.org/10.1016/j.isprsjprs.2014.05.011

Liggins, F., Vichi, A., Liu, W., Hogg, A., Kogou, S., Chen, J., & Liang, H. (2022). Hyperspectral imaging solutions for the non-invasive detection and automated mapping of copper trihydroxychlorides in ancient bronze. *Heritage Science*, *10*. https://doi.org/10.1186/s40494-022-00765-8

Lin, Y., Xu, C., & Lyu, S. (2019). Disease Regions Recognition on Mural Hyperspectral Images Combined by MNF and BP Neural Network. *Journal of Physics: Conference Series*, *1325*(1), 12095. https://doi.org/10.1088/1742-6596/1325/1/012095

Liu, G. L., & Kazarian, S. G. (2022). Recent advances and applications to cultural heritage using ATR-FTIR spectroscopy and ATR-FTIR spectroscopic imaging. *Analyst*, *147*(9), 1777–1797. https://doi.org/10.1039/D2AN00005A

Liu, L., Miteva, T., Delnevo, G., Mirri, S., Walter, P., de Viguerie, L., & Pouyet, E. (2023). Neural Networks for Hyperspectral Imaging of Historical Paintings: A Practical Review. *Sensors*, *23*(5). https://doi.org/10.3390/s23052419

Lloyd, S. P. (1982). Least Squares Quantization in PCM. In *IEEE TRANSACTIONS ON INFORMATION THEORY* (Vol. 28, Issue 2).

MacQueen, J. B. (1967). Some Methods for Classification and Analysis of MultiVariate Observations. In L. M. Le Cam & J. Neyman (Eds.), *Proc. of the fifth Berkeley Symposium*

on Mathematical Statistics and Probability (Vol. 1, pp. 281–297). University of California Press.

Mactaggart, P., & Mactaggart, A. (2007). *'White Pigments' In: Pigment ID using Polarised Light Microscopy*. https://academicprojects.co.uk/white-pigments/

Martinez, K., & Cupitt, J. (2005). VIPS - a highly tuned image processing software architecture. *IEEE International Conference on Image Processing 2005*, *2*, II–574. https://doi.org/10.1109/ICIP.2005.1530120

MATLAB. (n.d.). *HDF5 Files - MATLAB & Simulink - MathWorks United Kingdom*. Retrieved March 20, 2023, from https://uk.mathworks.com/help/matlab/hdf5-files.html

Maxwell, A. E., Warner, T. A., & Fang, F. (2018). Implementation of machine-learning classification in remote sensing: an applied review. *International Journal of Remote Sensing*, *39*(9), 2784–2817. https://doi.org/10.1080/01431161.2018.1433343

Mazzinghi, A., Ruberto, C., Castelli, L., Czelusniak, C., Giuntini, L., Mandò, P. A., & Taccetti, F. (2021). MA-XRF for the Characterisation of the Painting Materials and Technique of the Entombment of Christ by Rogier van der Weyden. *Applied Sciences*, *11*(13). https://doi.org/10.3390/app11136151

Mccarthy, B., & Giaccai, J. (2021). *Scientific Studies of Pigments in Chinese Paintings*. https://repository.si.edu/bitstream/handle/10088/111861/SSOPICP_Online_wCover.pdf?sequence=1

McCreery, R. L. (2000). *Raman Spectroscopy for Chemical Analysis*. John Wiley & Sons, Inc. https://doi.org/10.1002/0471721646

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133. https://doi.org/10.1007/BF02478259

McGillivray, M., & Duffy, C. (2017). New Light on the Sir Gawain and the Green Knight Manuscript: Multispectral Imaging and the Cotton Nero A.x. Illustrations. *Speculum*, *92*(S1), S110–S144. https://doi.org/10.1086/693361

McKinsey Global Institute. (2021). *The state of AI in 2021*. https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/McKinsey%20Analytics/Our%20Insights/Global%20survey%20The%20state%20of%20AI%20in%202021/Global-survey-The-state-of-AI-in-2021.pdf

McLaren, K., & Rigg, B. (1976). XII-The SDC Recommended Colour-difference Formula: Change to CIELAB. *Journal of the Society of Dyers and Colourists*, *92*(9), 337–338. https://doi.org/https://doi.org/10.1111/j.1478-4408.1976.tb03300.x

MegaVision. (2023). *MegaVision Cultural Heritage - EurekaVision*. https://www.mega-vision.com/cultural_heritage.html

Melit Devassy, B., & George, S. (2020). Dimensionality reduction and visualisation of hyperspectral ink data using t-SNE. *Forensic Science International*, *311*, 110194. https://doi.org/https://doi.org/10.1016/j.forsciint.2020.110194

Melit Devassy, B., George, S., & Nussbaum, P. (2020). Unsupervised Clustering of Hyperspectral Paper Data Using t-SNE. *Journal of Imaging*, *6*(5). https://doi.org/10.3390/jimaging6050029

Mellema, R. L., Sukir, & Hood, M. (1991). *Wayang puppets : carving, colouring, and symbolism*.

Miliani, C., Rosi, F., Brunetti, B. G., & Sgamellotti, A. (2010). In situ noninvasive study of artworks: The MOLAB multitechnique approach. *Accounts of Chemical Research*, *43*(6), 728–738. https://doi.org/10.1021/ar100010t

Mindat. (2023). *Galena from Peru*. https://www.mindat.org/locentries.php?p=5869&m=1641

Moros, J., Llorca, I., Cervera, M. L., Pastor, A., Garrigues, S., & de la Guardia, M. (2008). Chemometric determination of arsenic and lead in untreated powdered red paprika by diffuse reflectance near-infrared spectroscopy. *Analytica Chimica Acta*, *613*(2), 196–206. https://doi.org/10.1016/j.aca.2008.02.066

Mulholland, R., Howell, D., Beeby, A., Nicholson, C. E., & Domoney, K. (2017). Identifying eighteenth century pigments at the Bodleian library using in situ Raman spectroscopy, XRF and hyperspectral imaging. *Heritage Science*, *5*(1), 43. https://doi.org/10.1186/s40494-017-0157-y

Ngan, Q. (2018). The Significance of Azurite Blue in Two Ming Dynasty Birthday Portraits. *Metropolitan Museum Journal*, *53*, 48–65. https://doi.org/10.1086/701739

Nguyen, Q. H., Ly, H.-B., Ho, L. S., Al-Ansari, N., Le, H. Van, Tran, V. Q., Prakash, I., & Pham, B. T. (2021). Influence of Data Splitting on Performance of Machine Learning Models in Prediction of Shear Strength of Soil. *Mathematical Problems in Engineering*, *2021*, 4832864. https://doi.org/10.1155/2021/4832864

Orillaneda, B. C. (2016). *"Of Ships and Shipping: The Maritime Archaeology of Fifteenth Century CE Southeast Asia."* Springer.

perClass BV. (n.d.). *perClass Mira GUI for hyperspectral image analysis*. Retrieved March 20, 2023, from https://www.perclass.com/perclass-mira/product

Petroviciu, I., Teodorescu, I., Albu, F., Virgolici, M., Nagoda, E., & Medvedovici, A. (2019). Dyes and biological sources in nineteenth to twentieth century ethnographic textiles from

Transylvania, Romania. *Heritage Science*, *7*(1), 15. https://doi.org/10.1186/s40494-019-0255-0

Picollo, M., Cucci, C., Casini, A., & Stefani, L. (2020). Hyper-Spectral Imaging Technique in the Cultural Heritage Field: New Possible Scenarios. *Sensors*, *20*(10). https://doi.org/10.3390/s20102843

Plesters, J. (1966). Ultramarine Blue, Natural and Artificial. *Studies in Conservation*, *11*(2), 62–91. https://doi.org/10.2307/1505446

Polak, A., Kelman, T., Murray, P., Marshall, S., Stothard, D. J. M., Eastaugh, N., & Eastaugh, F. (2017). Hyperspectral imaging combined with data classification techniques as an aid for artwork authentication. *Journal of Cultural Heritage*, *26*, 1–11. https://doi.org/10.1016/J.CULHER.2017.01.013

Prediktera. (n.d.). *Prediktera – Hyperspectral Imaging Software*. Retrieved March 20, 2023, from https://prediktera.com/#software

Price, M. (2000). A Renaissance of Color: Particle Separation and Preparation of Azurite for Use in Oil Painting. *Leonardo*, *33*(4), 281–288. http://www.jstor.org/stable/1576902

Qi-yue, Y., Ting, Z., Ya-nan, H., Sheng-jie, H., Xuan, D., Li, H., & Chun-guang, X. (2020). From natural dye to herbal medicine: a systematic review of chemical constituents, pharmacological effects and clinical applications of indigo naturalis. *Chinese Medicine*, *15*(1), 127. https://doi.org/10.1186/s13020-020-00406-x

R Core Team. (2022). *R: A Language and Environment for Statistical Computing*. https://www.R-project.org/

Raggetti, L. (2021). *Traces of Ink: Experiences of Philology and Replication*. Brill. https://doi.org/https://doi.org/10.1163/9789004444805

Ricklefs, M. C. (1969). An Invesntory of the Javanse Manuscript Collection in the British Museum. *Bijdragen Tot de Taal-, Land- En Volkenkunde*, *125*(2), 241–262. http://www.jstor.org/stable/27861032

Ricklefs, M. C., Voorhoeve, P., & Gallop, A. (2014). *Indonesian Manuscripts in Great Britain*.

Riese, F. M., Keller, S., & Hinz, S. (2020). Supervised and Semi-Supervised Self-Organizing Maps for Regression and Classification Focusing on Hyperspectral Data. *Remote Sensing*, *12*(1). https://doi.org/10.3390/rs12010007

Rodarmel, C., & Shan, J. (2002). Principal Component Analysis for Hyperspectral Image Classification. *Information Systems*, *62*(2), 115–115. http://eo1.gsfc.nasa.gov].

Rojas, R. (1996). The Backpropagation Algorithm. In R. Rojas (Ed.), *Neural Networks: A Systematic Introduction* (pp. 149–182). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-61068-4_7

Romano, F. P., Caliri, C., Nicotra, P., Di Martino, S., Pappalardo, L., Rizzo, F., & Santos, H. C. (2017). Real-time elemental imaging of large dimension paintings with a novel mobile macro X-ray fluorescence (MA-XRF) scanning technique. *Journal of Analytical Atomic Spectrometry*, *32*(4), 773–781. https://doi.org/10.1039/c6ja00439c

RRUFF. (2023a). *Cinnabar R070532 - RRUFF Database: Raman, X-ray, Infrared, and Chemistry*. https://rruff.info/cinnabar/display=default/R070532

RRUFF. (2023b). *Orpiment R060105 - RRUFF Database: Raman, X-ray, Infrared, and Chemistry*. https://rruff.info/orpiment/display=default/R060105

Salouan, R., Safi, S., & Bouikhalene, B. (2014). A Comparison between the Self-Organizing Maps and the Support Vector Machines for Handwritten Latin Numerals Recognition. *International Journal of Innovation and Scientific Research*, *7*, 50–56.

Satopaa, V., Albrecht, J., Irwin, D., & Raghavan, B. (2011). Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior. *2011 31st International Conference on Distributed Computing Systems Workshops*, 166–171. https://doi.org/10.1109/ICDCSW.2011.20

Saverwyns, S., Currie, C., & Lamas-Delgado, E. (2018). Macro X-ray fluorescence scanning (MA-XRF) as tool in the authentication of paintings. *Microchemical Journal*, *137*, 139–147. https://doi.org/10.1016/j.microc.2017.10.008

Schafer, E. H. (1955). Orpiment and Realgar in Chinese Technology and Tradition. *Journal of the American Oriental Society*, *75*(2), 73–89. https://doi.org/10.2307/595009

Singh, M. (2020). Pigment Analysis of Palm Leaf Manuscripts of India. *Current Science*, *118*, 285–292. https://doi.org/10.18520/cs/v118/i2/285-292

Smith, M. L. (1999). "Indianization" from the Indian Point of View: Trade and Cultural Contacts with Southeast Asia in the Early First Millennium C.E. *Journal of the Economic and Social History of the Orient*, *42*(1), 1–26. http://www.jstor.org/stable/3632296

Solé, V. A., Papillon, E., Cotte, M., Walter, P., & Susini, J. (2007). A multiplatform code for the analysis of energy-dispersive X-ray fluorescence spectra. *Spectrochimica Acta Part B: Atomic Spectroscopy*, *62*(1), 63–68. https://doi.org/10.1016/J.SAB.2006.12.002

Specim. (n.d.). *Specim IQ Studio - Specim*. Retrieved March 20, 2023, from https://www.specim.com/iq/specim-iq-studio/

Spectral Evolution. (2023). *SR-3500 - Spectral Evolution*. https://spectralevolution.com/products/hardware/compact-lab-spectroradiometers/sr-3500/

Spring, M., Grout, R., & National Gallery, T. (2002). *The Blackening of Vermilion: An Analytical Study of the Process in Paintings*. www.nationalgallery.co.uk

Stanley, T. (2017). The Examination and Analysis of Dunhuang and Turfan manuscript materials at Princeton University Library's East Asian Library. *Journal of the American Institute for Conservation*, *56*(3–4), 194–210. https://doi.org/10.1080/01971360.2017.1392105

Stanley, T. (2018). Black and blue printing ink analysis by XRF, DRIFTS and Raman spectroscopy of recently discovered Gutenbergian Ars minor fragments. *Journal of the American Institute for Conservation*, *57*(4), 203–220. https://doi.org/10.1080/01971360.2018.1519359

Tarling, N. (1992). *The Cambridge History of Southeast Asia Volume one - From Early Times to c.1800* (N. Tarling, Ed.). Cambridge University Press.

Tarling, N. (2008). *THE CAMBRIDGE HISTORY OF SOUTHEAST ASIA VOLUME TWO The nineteenth and twentieth centuries* (N. Tarling, Ed.). Cambridge University Press.

Thanh Noi, P., & Kappas, M. (2017). Comparison of Random Forest, k-Nearest Neighbor, and Support Vector Machine Classifiers for Land Cover Classification Using Sentinel-2 Imagery. *Sensors (Basel, Switzerland)*, *18*(1). https://doi.org/10.3390/s18010018

ThermoFisher Scientific. (2019). *Niton XL3t GOLDD+ XRF Analyzer*. https://www.thermofisher.com/order/catalog/product/XL3TGOLDDPLUS

Uglietti, C., Gabrielli, P., Cooke, C. A., Vallelonga, P., & Thompson, L. G. (2015). Widespread pollution of the South American atmosphere predates the industrial revolution by 240 y. *Proceedings of the National Academy of Sciences*, *112*(8), 2349–2354. https://doi.org/10.1073/pnas.1421119112

Vahur, S., Teearu, A., Peets, P., Joosu, L., & Leito, I. (2016). ATR-FT-IR spectral collection of conservation materials in the extended region of 4000-80 cm–1. *Analytical and Bioanalytical Chemistry*, *408*(13), 3373–3379. https://doi.org/10.1007/s00216-016-9411-5

Van Grieken, R., & Eds. (2002). *Handbook of X-Ray Spectrometry, second edition.*

Wang, B. (2023). Trade in Artists' Materials in Eighteenth- and Nineteenth-Century Canton. *Studies in Conservation*, 1–10. https://doi.org/10.1080/00393630.2023.2174397

Wilkinson, R. J. (1935). Early Indian Influence in Malaysia. *Journal of the Malayan Branch and the Royal Asiatic Society*, *13*, 1–16.

Willbourn, E. S. (1925). A List of Minerals found in British Malaya together with a description of their Properties, Composition, Occurrences and Uses. *Journal of the Malayan Branch of the Royal Asiatic Society*, *3*(3 (95)), 57–100. http://www.jstor.org/stable/41560448

Wong, M., Abeysinghe, W., & Hung, C. C. (2019). A Massive Self-Organizing Map for Hyperspectral Image Classification. *Workshop on Hyperspectral Image and Signal Processing, Evolution in Remote Sensing*, *2019-September*. https://doi.org/10.1109/WHISPERS.2019.8921093

Yahya, F. (2021). Illustrated and Illuminated Manuscripts of the Dalāʾil al-khayrāt from Southeast Asia. *Journal of Islamic Manuscripts*, *12*(3–4), 529–581. https://doi.org/https://doi.org/10.1163/1878464X-01203012

Yuningsih, E., Matsueda, H., & Rosana, M. (2014). Epithermal Gold-Silver Deposits in Western Java, Indonesia: Gold-Silver Selenide-Telluride Mineralization. *Indonesian Journal on Geoscience*, *1*. https://doi.org/10.17014/ijog.v1i2.180

Zuena, M., Legnaioli, S., Campanella, B., Palleschi, V., Tomasin, P., Tufano, M. K., Modugno, F., La Nasa, J., & Nodari, L. (2020). Landing on the moon 50 years later: A multi-analytical investigation on Superficie Lunare (1969) by Giulio Turcato. *Microchemical Journal*, *157*, 105045. https://doi.org/https://doi.org/10.1016/j.microc.2020.105045

Appendix A

Flowcharts for Multiple Steps of Automated Clustering

Appendix A.1 – Flowchart for the Clustering of Large Collections – Data Reduction

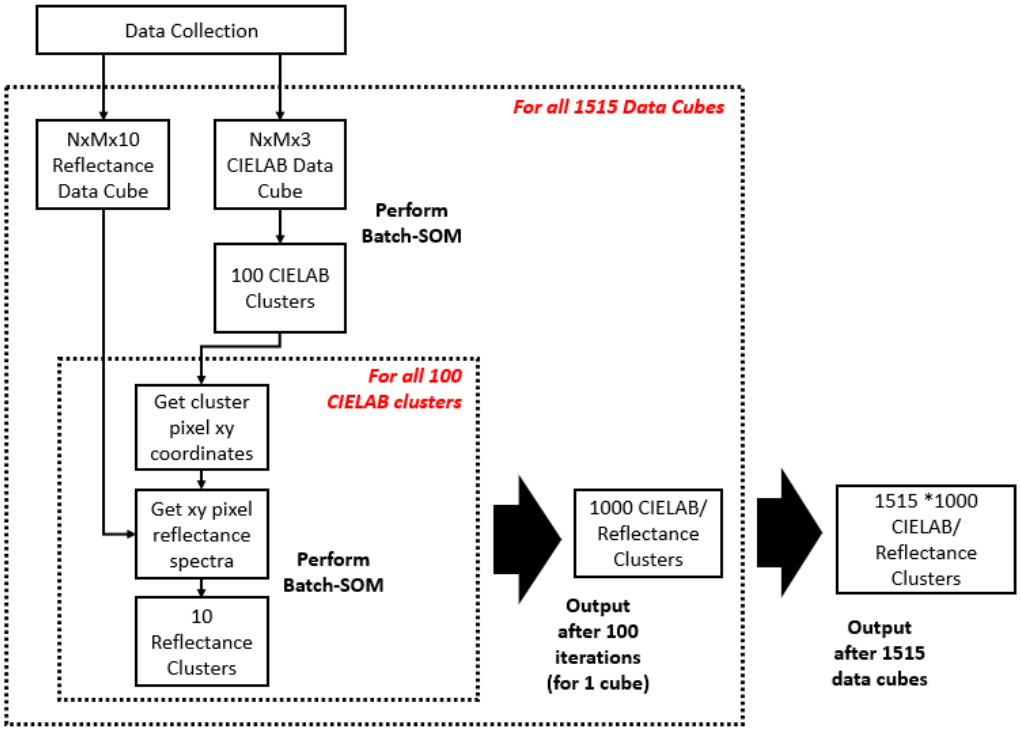

Appendix A.2 – Flowchart for the Clustering of Large Collections – Repeated Clustering

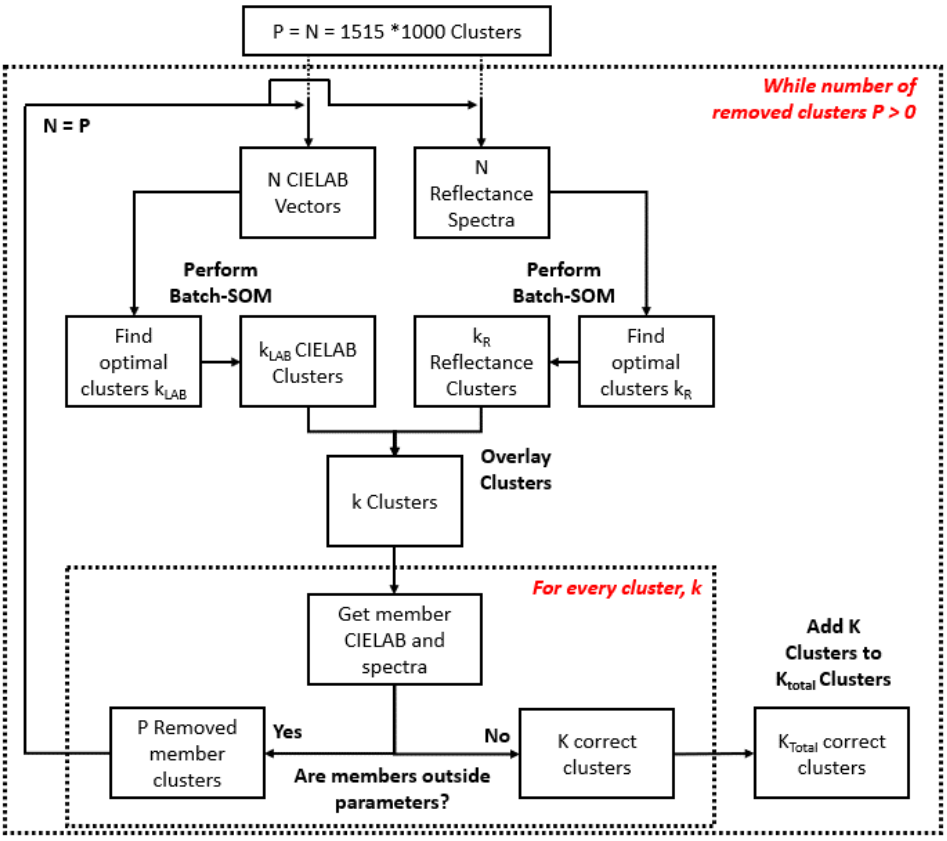

Appendix A.3 – Flowchart for the Clustering of Large Collections – Filtering

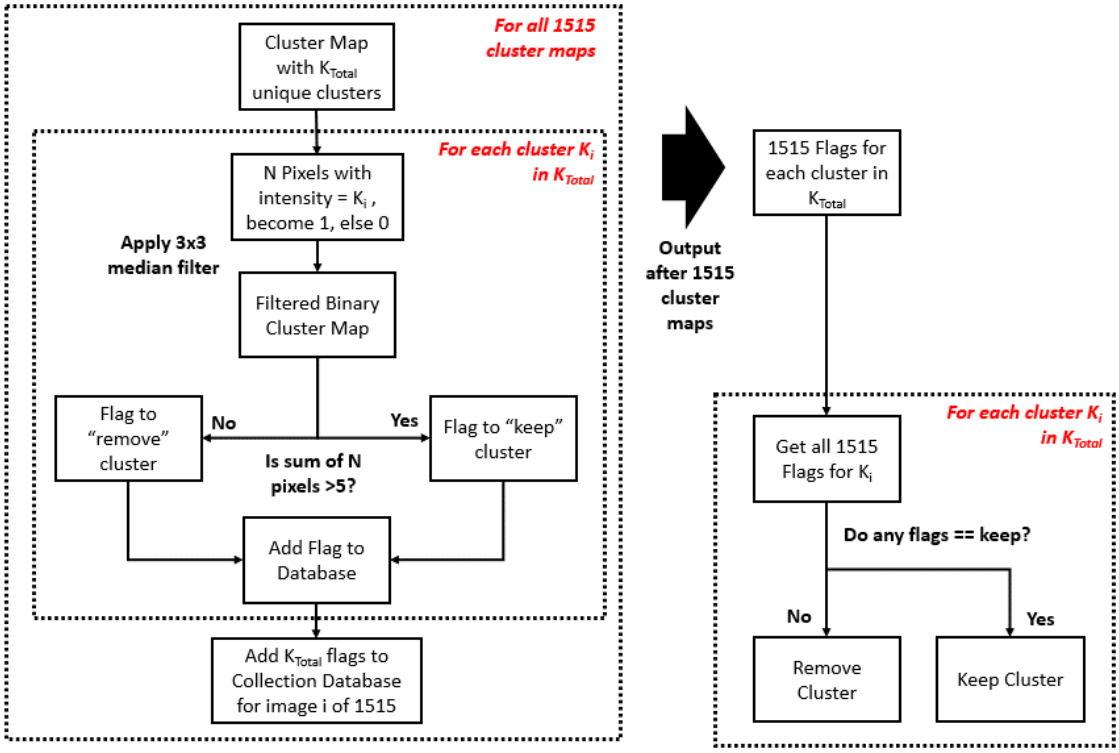

Appendix A.4 – Flowchart for the Clustering of Large Collections – Hierarchical Merging