Cognition, Technology & Work

Designing for Human-Agent collectives - Display considerations

--Manuscript Draft--

Manuscript Number:	CTWO-D-17-00008R1
Full Title:	Designing for Human-Agent collectives - Display considerations
Article Type:	S.I. : Doing more with fewer people
Keywords:	Human-Agent Teaming, Autonomy, Agent-based systems, Human-Machine Interface
Corresponding Author:	Dale Richards Coventry University Coventry, UNITED KINGDOM
Corresponding Author Secondary Information:	
Corresponding Author's Institution:	Coventry University
Corresponding Author's Secondary Institution:	
First Author:	Dale Richards, PhD
First Author Secondary Information:	
Order of Authors:	Dale Richards, PhD
	Alex Stedmon, PhD
Order of Authors Secondary Information:	
Funding Information:	GAMMA — Dr Dale Richards
Response to Reviewers:	Many thanks for all reviewer comments. I believe the vast majority have been addressed, while some of the others would require a new direction which we are considering for another paper.

Designing for Human-Agent collectives - Display considerations

Dale Richards[1] and Alex Stedmon
Coventry University, UK

The adoption of unmanned systems is growing at a steady rate, with the promise of improved task effectiveness and decreased costs associated with an increasing multitude of operations. The added flexibility that could potentially enable a single operator to control multiple unmanned platforms is thus viewed as a potential game-changer in terms of both cost and effectiveness. The use of advanced technologies that facilitate the control of multiple systems must lie within control frameworks that allow the delegation of authority between the human and the machine(s). Agent-based systems have been used across different domains in order to offer support to human operators, either as a form of decision support offered to the human or to directly carry out behaviours that lead to the achievement of a defined goal. This paper discusses the need for adopting a human-agent interaction paradigm in order to facilitate an effective human-agent partnership. An example of this is discussed, in which a single human operator may supervise and control multiple unmanned platforms within an emergency response scenario.

Key Words: Human-Agent Teaming, Autonomy, Agent-based systems, Human-Machine Interface

Introduction

Autonomy is a technological paradigm that allows the software of a socio-technical system to 'think for itself' and to conduct tasks in order to achieve goals set by a human operator (Baxter & Richards, 2010). There are different ways in which the use of *autonomy* can be introduced into a system, but for the purpose of this paper we will discuss the use of this technology as being represented by software agents.

Over the past decade there has been a considerable move from the traditional view of autonomy as being embodied within a physical robot that possesses some degree of intelligence, to a much wider perception of applications such as those demonstrated within distributed sensor networks, training decision support aids, intelligent buildings and also unmanned vehicles (see Grivault et al, 2016). Such complex agent-based systems tend to be situated within environments from which sensory inputs may be collected, whilst requiring little human direction in order to achieve the allocated goals (Jennings, Sycara & Wooldridge, 1998). Furthermore, Heath & Hill (2010) suggested that agent-based systems also tend to share properties with other agents in order to complete tasks, and thus we may perceive them as a collective whole (rather than the sum of individual agents - or even belonging to the same class of agents/platforms).

An agent may be defined as a self-contained entity that is able to interact with the environment and other agents (Weimer, Miller & Hill, 2016). An agent may also be viewed as a form of artificial intelligence (AI), and indeed these terms are often interchanged in some fields, such as robotics (Gigliotta & Nolfi, 2012). Early applications of software agents were portrayed as operating within

[1] Corresponding author: dale.richards@coventry.ac.uk

virtual environments (VEs), adapting dynamically with little, or no, human interaction (for example, Brooks, 1991; Beer, 1995; Clark, 1997). In these examples we begin to view agents as not individual component of a system, but rather as a network of data-hungry entities that not only gather information, but share and exchange it with other agents in order to further their understanding of their environment - and ultimately allow them to step closer to achieving their goal.

We can therefore view the introduction of intelligent software agents as existing across a network, providing support to the human in terms of supporting critical decision-making while also providing a means by which some of the more routine or mundane tasks may be delegated to the agent. This would then ‘free up’ the human operator's cognitive resources, allowing them to focus on the overall global objective/s. This act of delegated autonomy would therefore create a partnership between the human and agent that would allow both elements to focus on the strengths they bring to the partnership. For example, agents should be better suited for tasks that involved laborious, mundane and/or complex tasks to be carried-out, while the human operator is better at dealing with dynamic and uncertain factors.

One of the many domains where we are witnessing more use of agent-based autonomy is within the application of unmanned vehicles (Richards & Stedmon, 2016). The advances in this technology are not only focused on the improvement of sensors utilised by the vehicle, but also includes the use of complex automation to facilitate command and control processes that can assist in improving task effectiveness. The introduction of more advanced automation (or the delegation of decision making to the system via autonomy) to those platforms provides the human within the system with a wider capability and flexibility in achieving their goals.

This paper discusses the design of a display that allows for the human operator to define goals that may then be delegated to the agent-based system to conduct. The nature of human-agent interaction will be discussed and conveyed in relation to an example of how a single operator can task multiple unmanned systems.

The Human-Agent Partnership

The ability for a single human operator to delegate manual control to an agent-based system was explored under the Human-Agent Mission Planning for Emergency Responders (HAMPER) project (Richards et al., 2014). This system not only allows the operator to issue direct commands to individual assets under their control, but to effectively delegate control to an agent-based system that generates its own decisions that allow the agents to achieve a designated goal.

It is important that the delegation of control between the human and the agent is safe, reliable and effective but it must also take into account the human perspective of interaction by ensuring the system is usable and promotes confidence in what the system is doing (Richards et al., 2014). The use of highly complex systems to achieve goals has resulted in the need for new paradigms in terms of human supervisory control and a better understanding the nature of cognition between the human and the technology (Richards & Stedmon, 2016).

There are many ways of defining models of co-operation between intelligent software agents; ranging from Joint Intentions (Levesque, Cohen & Nunes, 1990) to Shared Plans (Grosz & Sidner, 1990), and beliefs (Bratman, 1984). These different approaches provide us with an important aspect of understanding exactly how an agent behaves within given states and contexts: in essence how the agent draws on information within its environment and then acts upon it by making a decision. This process, of course, holds at its core a very human issue that will determine whether the agents are performing in an appropriate manner, or whether the goal is attainable. When we examine human team performance, the role and importance of *trust* is a significant factor in determining the efficiency of that team and likelihood of achieving their goal (Richards, 2017). There are have been many different attempts to define the concept of trust, ranging from psychological, sociological and philosophical concepts being related to: predictability (Gambetta, 1988); a willingness to allow action to occur (Mayer et al, 1995); or a set of interrelated factors (McKnight & Chervany, 2001). However, a general consensus tends to be that it is a complex construct that not only has an underlying psychological component, but also a process-driven causal link that serves to reinforce attitudes and actions across the overall system.

In order for the human to trust the system, a process of delegated decision making needs to be conveyed back to the human operating the system in order to understand not only what the system is doing, but what the system plans to do next. We may therefore view the agent element of the system alongside the human, sharing the same goals, and in some instances sharing very similar conceptual plans and expectations. Johnson-Laird (1983) proposed that it is possible to conceptualise human perceptions and knowledge in terms of structured mental representations – mental models, and that these models may form the basis of not just mental representations for objects and processes but also how humans reason (Johnson-Laird & Byrne, 1991). However, the nature of these mental representations may not be as well defined as we would wish and thus present fuzzy constructs (Gray, Zanre, & Gray, 2014) that still act in ways that possess meaningful relationships between concepts within the model (Novak & Canas, 2008). However, it is also worth considering the nature of how such cognitive elements are not mere abstractions that exist within the confines of our brain, but are influenced by external forces within the environment, or indeed other psycho-sensory mechanisms (Chemero, 2009). This is the nature of embodied cognition, whereby models of mental representation are viewed in relation to wider influencing factors, rather than as isolated cognitive models: the cognition of a system is therefore situated within a wider context of use. However we might model perceived cognition between the human and the agent, we can suggest that the nature of human and agent representation of achieving goals exists in a very similar manner; in that agents may also need to operate with fuzzy conceptual ideas of the external world that may be *shaped* by external influences.

For the purpose of the HAMPER project the agent-based system was based on fuzzy logic. This approach has been used across several different fields and has been applied to intelligent unmanned systems (e.g. Cetin et al, 2011; Pearson et al, 2014; Zeng et al, 2013; Ernest & Cohen, 2015). These systems are often defined as operating within a level of bounded rationality and thus may arrive at decisions that require only partial truths, rather than binary decisions (i.e. true or false). The concept of 'bounded rationality' was first proposed to help explain economic systems, stressing that all agents within that system are complex and possess imperfect knowledge whilst also likely to display limited (and differing) abilities (Simon, 1955). Key to the nature of these *imperfect agents* is the role of communication between themselves and others. This is key when we approach decision-making and planning activities within a human-agent system; as there has to be at least two representations of a given decision, or series of tasks, that can be represented from either a human or agent-based perspective. In order to achieve trust within an agent-based system a degree of understanding how the agent model is represented is important (as the agent will display different behaviours in order to achieve a goal). We can thus view this (in its simplest way) as two models of representation: the human mental model and the agent system model (Figure 1).

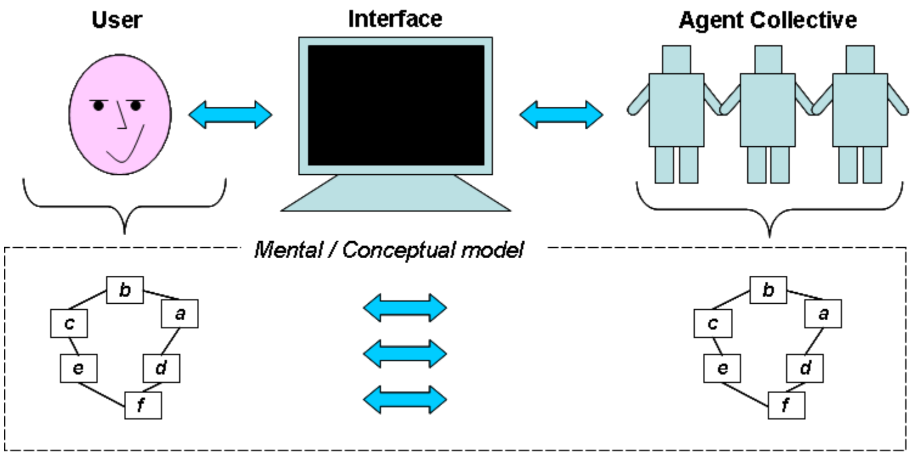

Figure 1 – Human-Agent conceptual representation of tasks

In order to manage and allow the human to interact with the agent-based system it is essential to have a notional framework of how the human delegates tasks to the agent. This may either be discrete single tasks that may require rule-based (bounded) responses from the agent, or allow the agent to have more freedom in which to attain the goal. Richards & Stedmon (2016) suggest that several frameworks exist that have been used to control/enable autonomous systems in different domains; e.g. ranging from unmanned systems (PACT)[2] and automotive applications (NHTSA)[3].

This presents a challenge in terms of how information is presented to the human, as this directly influences the nature of human-agent interaction. Information presented via the Human Machine Interface (HMI) needs to not only present the relevant information to the user in terms of assisting in the building and maintaining of their Situation Awareness (SA) but also for the agent system to update the human in terms of what the agents are actually doing and might be doing in the near future (i.e. to support user expectations).

The HAMPER Human Machine Interface

At the core of the HAMPER project was the requirement to build a display that not only conveyed information to the user in terms of mission status and progress, but allowed the human to directly interact with the agent software that was responsible for controlling a range of aerial and ground-based unmanned vehicles. This was achieved via an interface where a single operator could delegate authority of control to different agents that, in turn, would govern other unmanned vehicles. The HAMPER system would also allow the human to use the system without software agent support, but these elements are not considered relevant in discussing the nature of human-agent interaction within the HAMPER system.

---

[2]Pilot Authorisation and Control of Tasks. See Bonner, Taylor & Miller (2000).
[3]National Highway Traffic Safety Administration

How information is displayed on the HMI is of vital importance in terms of an effective means by which the user can develop and maintain good SA while also providing the ability for the user to interact with the system under their control. In order to achieved this there are a number of good HMI design principles that can be used as guidance (such as scenario-based design or human-centred design) but this paper considers more the nature of human-agent interaction and how information should be focussed on the nature of information that is required to facilitate human-agent interaction.

The HMI developed within the HAMPER project allows the human to manually control unmanned assets (through the use of automated scripts that respond to waypoint functions) or to switch to an autonomous mode (whereby the system decides which assets are better situated to achieve the goal). This requires a HMI that provides flexibility in terms of not only providing direct input commands to the assets but also the ability to set goals that can be reflected by different agent behaviours as tasks are distributed amongst the agents in a dynamic operational environment (see Richards et al, 2014).

The primary display for the HAMPER operator consists of multiple viewing panes, of which some may be configured by the operator (see Figure 2).

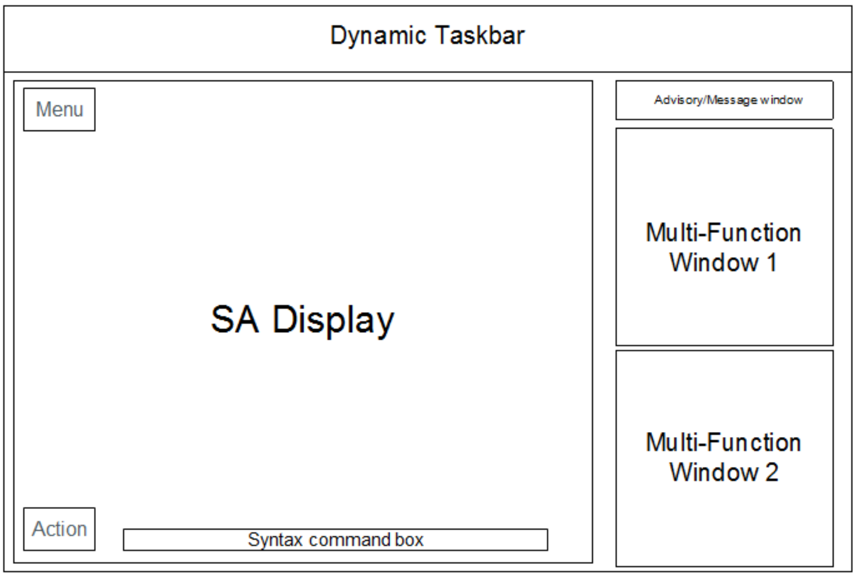

Figure 2 - HAMPER Human-Machine Interface

The rationale behind this initial display layout is based on *Gestalt* principles of symmetry and form, whilst using location as a key perceptual cue. This allows the user to ascertain which part of the display is related to certain functions or information. At the top of the display the dynamic taskbar consists of several information elements that either reflect asset state or provide the ability for the operator to populate the multi-function windows (MFW) with different views (e.g. imagery, messages, tasks). To the left of this aspect of the display are the available unmanned assets (as represented in Figure 3). In this example, we can see that two UAVs and two unmanned ground vehicles (UGVs) are available for tasking. In this sense we refer to this aspect of the system as the unmanned autonomous system (UAS).

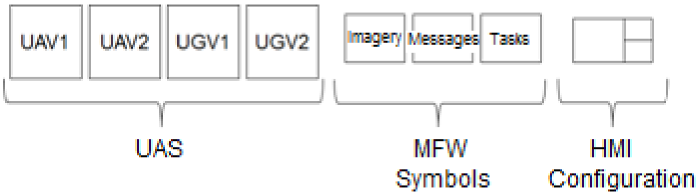

Figure 3 - Dynamic Taskbar

On the far right of this display the operator can select a widescreen mode of the SA display, allowing the map view to fill the screen when the multi-function windows are not required for operational purposes. However, during HAMPER missions users tend to refer back to the MFWs and keep them displayed more often than not.

The platform element of the taskbar allows the operator to view state information relating to each of the assets and also the ability to select the asset for issuing commands. Each UAS that is operational displays standard information relating to: call sign, flight data, behaviour mode (e.g. Waypoint following), data link, and available sensors (see Figure 4).

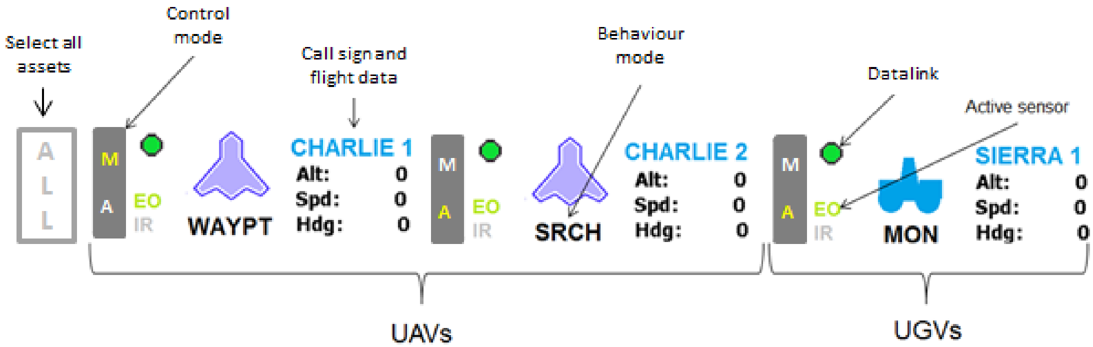

Figure 4 – Dynamic Taskbar for UAS

In order to assist the operator in tasking the UAS a message box is incorporated within the HAMPER HMI. This allows the operator the ability to observe information associated with each of the assets, tracks or points of interest on the SA display (see Figure 5).

Figure 5 – Dialogue order box (used for conveying information)

In this example the operator has selected a point of interest (POI) on the SA display that refers to a location where they have instructed one of the UAVs (Charlie 2) to perform an associated behaviour (i.e. Monitor). This allows the operator to maintain good SA over the mission tasks and the allocation of functions across available assets. This is also linked to mission goal progress.

The Syntax Command Box is displayed on the bottom of the SA display and is used for building/issuing commands to the UAS. This represents an important aspect of the HAMPER HMI as it provides the means by which the user interacts with the agents. The nature of the commands needed within the dialogue box has a formal structure that will prompt the user to populate the appropriate information required to complete a command string. This in effect creates a syntax that the user must adopt in order to not only issue commands to the agents, but which also allows the user to check the elements within any order that is being constructed (before it is sent to the agents).

As soon as the bottom left button on the SA display is pressed (ACTION) the syntax box will change to the command format that will begin to provide prompts to the operator. In order to send a command to the UAS, the operator will always have to adopt a constant command string that will be constructed of command elements. Thus, if the operator wished to command one of the UAVs to search at a particular location they would have to fulfil the following semantic command string: ASSET>ACTION>LOCATION. In order to assist the operator a prompt is issued within the syntax box in order to highlight which part of the string would be required to complete the command. At the end of a string a ACCEPT or REJECT button provide the options for the operator to either execute the agent behaviour or revert back to a manual task (see Figure 6).

Figure 6 – Syntax order box (used for issuing commands)

The construction of orders in this manner provides an intuitive way for tasking single assets (or *singletons*) under the operator's command. However, in some instances the operator may wish to ask for decision support and simply build the command string and define the asset as ALL. This will prompt the HAMPER system to examine all assets within the command infrastructure and issue commands that would reflect an optimised and efficient response to the goal as defined by the operator. A task assigned to all assets within HAMPER may reflect mission critical tasks, such as searching for something of interest and where a return of assets is urgently required.

Agent Behaviours

When Autonomy mode is engaged within the HAMPER system this allows the operator to use intelligent agent decision support. In direct comparison to Manual mode, this instance of HAMPER provides decision support to the operator by presenting tasks as completed plans rather than the operator having to construct native commands that may be built into a behaviour (or task). This mode also allows the operator to command multiple assets and perform tasks at a group level; with the agent software taking into account the available resources that are available to complete the task and the relative constraints associated with the environment/task.

There are several commands that will allow the operator to select specific tasks that facilitate intelligent agent behaviours. While these individual behaviours are described in Table 1 it is assumed that when in Autonomous mode the agents will possess primitive intelligence that will maintain safe separation with other air/ground agents and associated static zones/objects. There is an assumption that the UAS has sensor abilities that are capable of such integrity (although this was considered out of scope for HAMPER).

BEHAVIOUR	DESCRIPTION	EFFECT
LOIT (Loiter)	Initiates LOITER mode.	Dynamic taskbar will indicate LOITER mode, and task window will indicate this also.
RCVR (Recovery)	Initiates UASs to return to base.	Dynamic taskbar will indicate RECOVERY mode, and task window will indicate this also.
WYPT (Waypoint)	Inserts Waypoint	Initiates and adds waypoint on map.
OBSV (Observe)	Initiates OBSERVE mode.	Dynamic taskbar will indicate OBSERVE mode, and task window will indicate this also.
SRCH (Search)	Initiates SEARCH mode.	Dynamic taskbar will indicate SEARCH mode, and task window will indicate this also.

Table 1 – Agent Behaviours

Example Behaviour - Search

In a typical search operation, agents might be commanded to perform a search in order to provide visual information back to the operator. In order to maximise the benefit and efficiency, agents must search in an optimised way in terms of synchronising with other manned or unmanned assets involved in the exercise. In order for this to occur it is necessary for all the agents within the system to be able to respond dynamically to other agents, the dynamically changing environment, and human inputs from the search team or operator.

Within the HAMPER system, the search behaviour is written specifically for UAVs while UGVs being controlled by the agents were not involved in search activities due to limitations of the UGV to rapidly deploy and search. Using the HAMPER system it is possible to perform a manual search by simply flying a UAV into a search area and instructing it to observe different points within that area. This is of course time intensive and places a high degree of cognitive effort on the human operator.

In order for the operator to engage an autonomous behaviour for the search function, they must first designate an area. This is done by either associating a search area with a POI, or defining and then dropping a search box onto the map within the SA display. In each of these instances it is required that the search area (i.e. a centre point and search radius) is displayed on the SA display and the search status of associated assets is updated on the dynamic task bar.

Different search behaviours may be achieved depending on which specific assets are available to the operator at the time of tasking. The composition of this 'team' could be constructed of single or multiple agents and potentially include humans on the ground or in the air. The key focus for the HAMPER system was the nature of how a single operator could direct a team of both agents and humans to share a common goal (in this instance to search a designated area together and to collaborate their activity in order to maximise effect). Figure 7 provides an overview of such teaming options, dependent on number of agents, number of humans, completion of search sectors, and updates to ‘close out’ search sectors (i.e. identify those not to be searched where intelligence updates have been used to narrow the continuing search in real time).

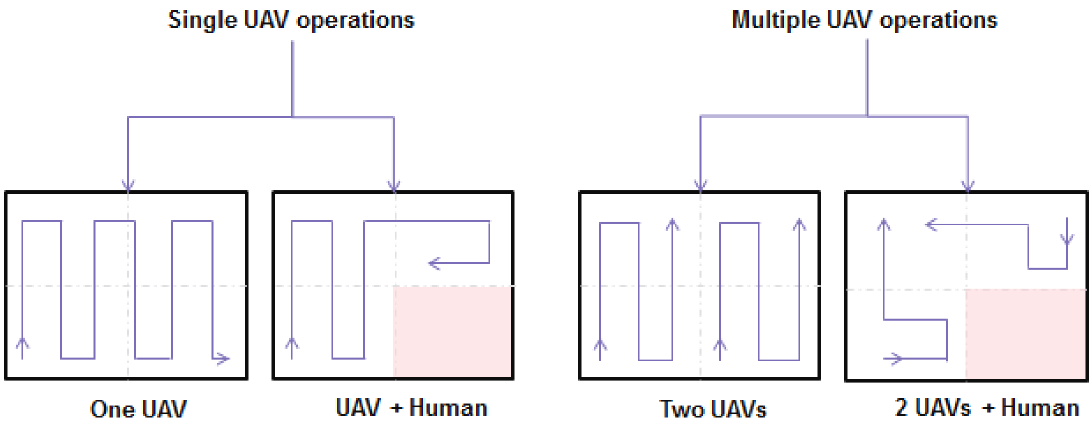

Figure 7 – Search behaviours

It was assumed that a human search party would start in a known area and remain in a sector of that search area. During the search activity this sector would update after a period of time by means of a message sent to the HAMPER system that would update all the agents.

At the beginning of the search (and entry into search box) it is assumed that no information is known regarding which part of the search box is more likely to contain the target. However, on entry into the search box the asset tasked to search it would adopt a pre-planned route as outlined above and indicated in Figure 8. On crossing the search box boundary (Figure 8 - b) the UAV will change status from ‘Ingress’ to ‘SEARCH’.

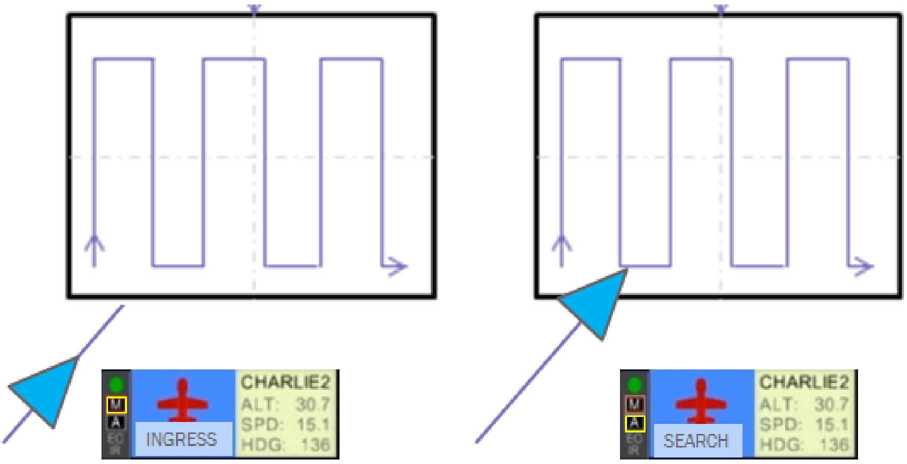

(a) UAV flying to Search Box (b) UAV entering Search Box

Figure 8 – Joining Search Box

It is important to allow the operator to observe the UAVs while in the search box undertaking autonomous search behaviour. Feedback is provided to the operator not only on the flight path that the UAV is following, but also the sensor footprint that the UAV is covering during the search. By displaying this it draws a zone around the UAV that will allow it to scan it's field of regard, as seen in Figure 9. The HAMPER HMI allows the operator to select this level of detail on the SA display if they so desire.

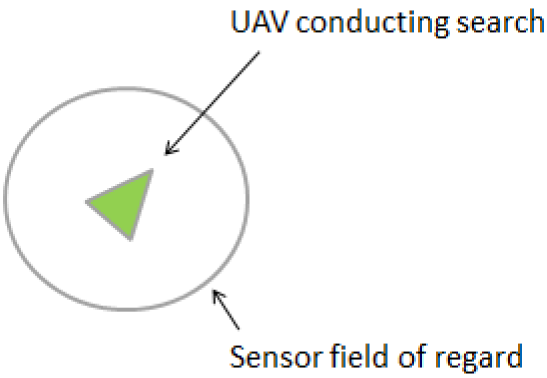

Figure 9 – UAV and sensor field of regard

As the UAV enters the search box and activates the planner, this dictates search flight path and a number of visit areas identified along the path. As each of these areas are 'visited' they may be visually changed within the search box in order to present feedback to the operator in terms of progress. This is shown in Figure 10. This is a function that is automatically de-cluttered on start-up and the operator may choose to populate the SA display with this information if they so desire.

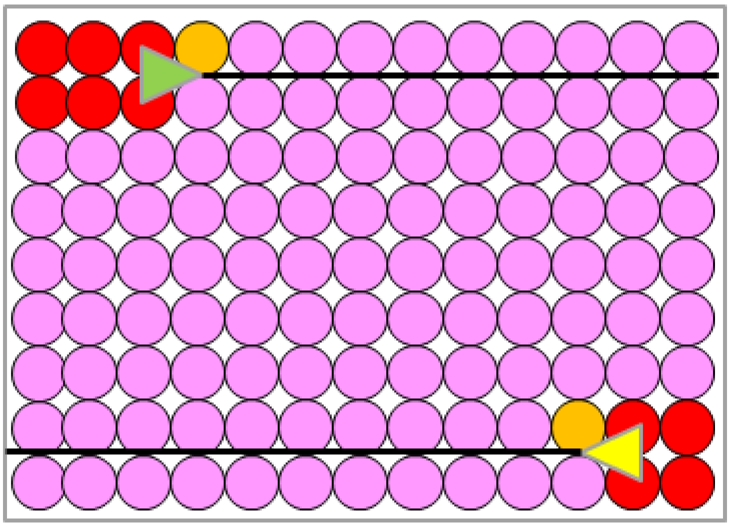

Figure 10 – Two UAVs within search box with searched areas completed (in red), current search area (orange) and areas to be completed (purple)

Multiple UAVs may be added to the task at any point of the search. Thus, an operator may choose to exploit new resources as they become available from other tasks either to increase the efficiency of the search or to complement the activities of existing agents in the system. For example, in Figure 11 below we can witness two UAVs assisting with a search mission that also includes a human search team.

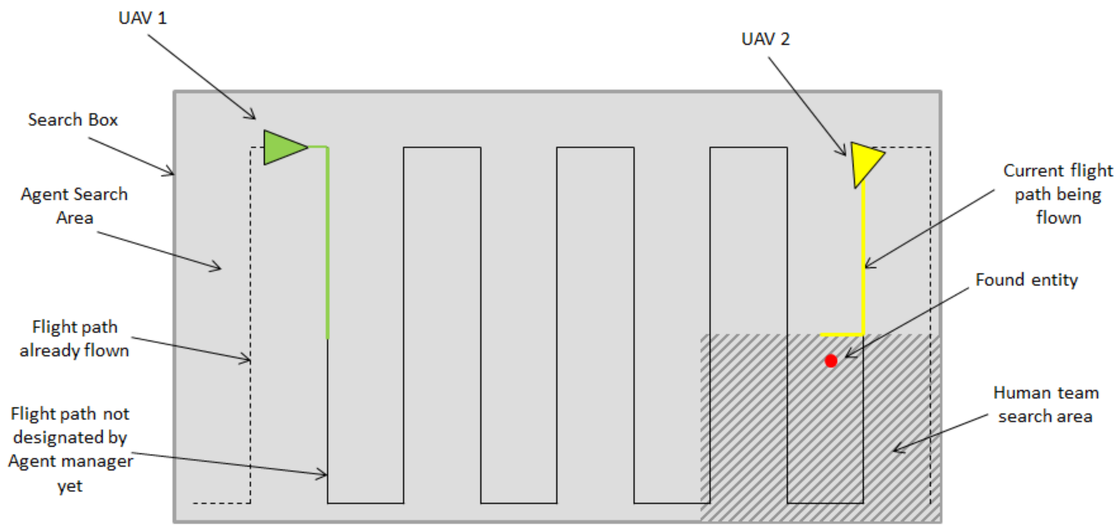

Figure 11 – Search Behaviour between human and agent actors

In this example we can see that a search box has been created and that two autonomous UAVs are searching the area. The greyed hatching represents an area that has been designated to a human search team and we can see the planned flight path for one of the assets has changed in order to accommodate the availability of the human team. On finding the ‘target’ (i.e. lost person or entity) the unmanned agents in the system will then execute the next behaviour associated with this process (e.g. one of the assets might loiter over the area and provide observation, while another returns to base or adopts another task behaviour).

The completed HAMPER HMI provides the human operator with the means to command and supervise multiple autonomous assets. By adopting a cognitive engineering approach to designing the HMI the operator's needs are of primary importance, allowing them to maintain a high level of SA in relation to the mission and also viewing and developing a deeper level of understanding about agent behaviour within the system associated with the unmanned systems. The completed HMI is shown in Figure 12.

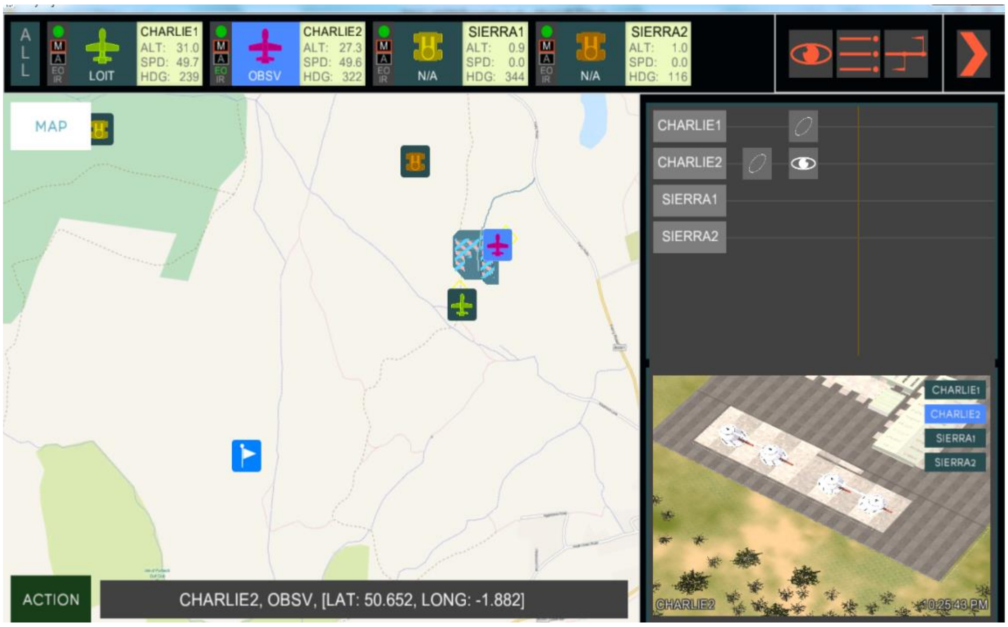

Figure 12 – HAMPER HMI during simulated mission scenario (showing full motion video and task displays

Discussion

In this context an autonomous agent is viewed as being able to operate without the direct intervention of a human, and possess a property that suggests they have control over their own actions and internal state (Castelfranchi, 1995). When we try to discuss the nature of autonomous agent behaviours we can often find ourselves reporting the behaviour in an anthropomorphic manner, almost akin to Folk Psychology, where we assign human attributes to the agent's state, beliefs and goals. However, with this approach to agent-based autonomy is an underlining assumption that the agents are self-reliant and are destined to continue their personal odyssey within the system/environment. This deterministic view is better suited to the earlier incarnations of autonomous agent systems (see Brooks, 1991; Beer, 1995; Clark, 1997) and not the emerging pattern we are now seeing whereby there is a human-agent partnership with both parties changing role and level of authority.

The HAMPER system as discussed in this paper provides such an example, with the mission goal being shared between the human and the agent; and at times the boundary between the two is less distinct in terms of task or indeed following orders. To a large extent the HAMPER system allows us to observe the nature of teaming between the human and agent(s); where all team members collaborate towards a common (and shared goal). According to Endsley (1995) in order for team members to be aware of each other, more information is required for effective team coordination.

Thus we can view the HAMPER system as representing the means by which the human builds their own SA of the mission situation as defined in HAMPER, but equally importantly provides cues back to the agent-based system as to the context of both the mission and the human operators engaged with the system (i.e. the human team on the ground and the intent of the human operator). Endsley & Jones (1997) suggest that sharing of goals between team members assists in building what they refer to as 'shared SA' that allows for a better understanding of intent between all members. When this level of shared understanding is absent or at a low level it has been found to have a negative impact on achieving the team goal (Bolstad & Cuevas, 2010). Previous studies have stated the importance of designing an appropriate HMI that allows for improved SA across team members (Biehl et al., 2007; Parush & Ma, 2012).

As noted earlier, trust is a critical component in the design of human-agent interaction, and the HMI presents the fundamental focal point that can determine whether this interaction can build an efficient interaction that can achieve a goal or perhaps create a source of distrust for the human. Hardin (2002) suggests that trust between two parties is built upon an expectation that actions will be pursued to fulfil self-interest and somewhere between this both parties will benefit. This 'embedded-interest' implies that behaviours are self-regulated and act towards individual goals that may either be shared or part of a wider achievement that leads to a shared goal. Key to Hardin's proposal is the assumption that we assume the other party to be rational and for them to make rational decisions. Therefore, in order to perceive whether a behaviour or decision is indeed rational we must present sufficient information to the other party in order for them to judge whether it is indeed rational. The HMI, if designed correctly to support user needs, may be viewed as the medium by which this relationship may flourish. Alternatively, if it fails to support the user's understanding of the system it may extinguish any sense of trust or shared understandings across the agents.

The nature of how we view shared cognition between human and agent team members is worth considering in relation to the HAMPER system. The radical movement on embodied cognition would argue that the use of mental models is not best placed to understand how components pertaining to human-agent cognition should be designed. However, Wilson (2002) reminds us that many cognitive activities may take place without a task or embodied context (such as daydreaming) and that adopting mental representations are still better suited to some cognitive activities (such as spatial reasoning and problem-solving). So, which approach may best suit the design of the visual component of the HAMPER system? It is possible to hold both views when considering human-agent interaction deign. One thing is clear when it comes to what we need to build into a HMI; in order to design the HMI we must extract interaction requirements from the human. Many techniques may be applied to achieve this, but the key aspect relies upon the mental construct and perception that the human holds when thinking about a system (or interaction) that has yet to take place in the physical world. The evaluation of an existing system may afford a better understanding under the guidance of embodied cognition; thus allowing the role that perceptual processing may have in determining our *situated cognition*. This is an important component when we consider the significance of tasks within the HAMPER system and the way in which they define both human and agent behaviours.

The use of graphical user interfaces is thought to reduce the complexity of the operating system and make it easier for human operators to use them and learn about underlying task-related principles (Wiedenbeck, 1999). As a result, any new format of interface that represents the point of contact for a new system is crucial in terms of the usability and indeed *learnability* of such a novel system. It is important that we view the interaction between the human and agent as a partnership in every sense of the term, with control passing seamlessly between the two. In many ways this partnership is no different to human-human partnerships that we seem in traditional team settings, with the success being based on a transparent and understandable communication between all parties involved in the interaction (Richards, 2017).

The use of autonomous software agents within UAVs presents an opportunity to maximise the chance of mission effectiveness and success. It is essential that the HMI not only allows the operator to issue low level commands (e.g. move to location X) but also to task abstract goals to the unmanned systems available (e.g. search area X and find Y). Thus, the HMI needs to allow both forms of interaction whilst also providing appropriate levels of information back to the operator in a timely manner. This can be made more complex when we consider the introduction of human agents within the scenario. For example, if we have set a goal for the agents to find something within a given search area, we can predict that the agents will collaborate and coordinate a plan together. However, if a human agent is inserted into this scenario the agents will have to be prompted by the human as to where the human agents are located (and what may be expected of them under such circumstances). This human-agent collective is demonstrated in the HAMPER system whereby a search area is defined and the agent system designates available assets to begin the search. When a human team is also located within this area, the agent system informs the existing agents where the human team is searching and accordingly adapts its behaviour accordingly. This allows for an optimisation of the search plan to ensure that available assets are used to maximum effect. For example, it would be a waste of resources for an agent to repeat a search activity in an area that another agent or human team is already searching within (quite apart from the inherent safety concerns of having disorganised and uncoordinated agents operating in the same area or close proximity to humans that may be unaware of their whereabouts or task requirements).

Conclusion

The key to achieving mission effectiveness is to provide the decision-maker with good tactical SA. In emergency situations uncertainty is the most prominent factor that can influence the resolution of a disaster and facilitate a state of normality that resembles events before the event occurred. Team members must not only make sense of what this means but also be able to achieve a collective shared awareness of what the reality of the situation is; and also how best to resolve (or contain) the consequences of the event using all the resources they have available. While agent systems can be of benefit in carrying out complex, routine, or dangerous tasks, there is significant benefit in examining the principles required for a human-agent collective team.

Acknowledgements

This work was funded under the Growing Autonomous Mission Management Applications (GAMMA) programme. The author would like to express his thanks to the several reviewers who provided valuable suggestions on an earlier draft of this paper. Similarly, many thanks to Dr John Huddlestone, Guest Co-Editor of CTW, for his *gentle* reminders and advice.

References

Baxter, J. W. & Richards, D. (2010) Whose goal is it anyway? User interaction in an autonomous system. In proceedings of the workshop on Goal Directed Autonomy, AAAI2010, Atlanta, 2010

Bonner, M.C., Taylor, R.M., Fletcher, K., and Miller, C. (2000). Adaptive automation and decision aiding in the military fast jet domain. In: CD Proceedings of the Conference on human Performance, Situation Awareness and Automation: User-Centered Design for the New Millennium (154-159). Savannah, GA.

Bratman, M. (1984) Two faces of intention. The Philosophical Review, 93(3), 375-405.

Brooks, R. A. (1991) Intelligence without representation. Artificial intelligence, 47(13), 139-159.

Beer, R.D. (1995) A dynamical systems perspective on agent-environment interaction. Artificial Intelligence, 72, 173-215.

Castelfranchi, C. (1995) Guarantees for autonomy in cognitive agent architecture. In M. Wooldridge and N. R. Jennings (eds.) Intelligent Agents: Theories, Architectures, and Languages (LNAI Volume 890), pp.56-70, Springer-Verlag.

Cetin, O., Kurnaz, S., & Kaynak, O. (2011) Fuzzy logic based approach to design of autonomous landing system for unmanned aerial vehicles. Journal of Intelligent & Robotic Systems, 61(1), 239-250.

Chemero, A. (2009). Radical Embodied Cognitive Science. Cambridge, MA: MIT Press.

Clark, A. (1997) Being there: Putting brain, body and world together again. Cambridge, MA: MIT Press.

Ernest, N. & Cohen, K. (2015) Fuzzy logic based intelligent agents for unmanned combat aerial vehicle control. Journal of Defence Management, 6: 139. doi:10.4172/2167-0374.1000139

Gigliotta, O. & Nolfi, S. (2012) On the coupling between agent internal and agent/environmental dynamics: development of spatial representations in evolving autonomous robots. Adaptive Behavior, 16, 148-165.

Gray, S.A., Zanre, E., & Gray, S.R.J (2014) Fuzzy cognitive maps as representations of mental models and group beliefs. Fuzzy Cognitive Maps for Applied Sciences and Engineering, 29-48.

Grivault,L., Fallah-Seghrouchni, A., Girard-Claudon, R. (2016) Agent-Based Architecture for Multi-sensors System Deployed on Airborne Platform. IEEE Internnational Conference on Agents (ICA), 28-30 Sept, 2016.

Grosz, B.J. & Sidner, C. (1990) Plans for Discourse. In Cohen, Morgan and Pollack (eds.), Intentions in Communications, MIT Press, Chapter 20, pp. 417-444.

Hardin, R. (2002) Trust and trustworthiness. New York: Russell Sage Foundation.

Heath, B. L. and R. R. Hill (2010) Some Insights into the Emergence of Agent-Based Modelling. Journal of Simulation 4(3):163-169.

Jennings, N. R., K. Sycara, and M. Wooldridge, (1998) A Roadmap of Agent Research and Development. Autonomous Agents and Multi-Agent Systems 1:7-38.

Johnson-Laird, P.N. (1983) Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness. Cambridge: Cambridge University Press.

Johnson-Laird, P. N. & Byrne, R. M. J. (1991) Deduction. Hillsdale, N. J.: Lawrence Erlbaum.

Levesque, H.J., Cohen, P.R., & Nunes, J.H.T. (1990) On acting together. Eighth National Conference on Artificial Intelligence, 94-99.

Novak, J.D. & Canas, A. J. (2008) The theory underlying concept maps and how to construct and use them. Technical Report IHMC CmapTools 2006-01 Rev 01-2008. Institute for Human and Machine Cognition. Florida, pp.36.

Pearson, D., An, E., Dhanak, M., von Ellenreider, K., & Beaujean, P. (2014) High-level fuzzy logic guidance system for an unmanned surface vehicle (USV) tasked to perform autonomous launch and recovery (ALR) of an autonomous underwater vehicle (AUV). *2014 IEEE/OES Autonomous Underwater Vehicles (AUV)*, pp. 1-15, Oct 2014.

Richards, D. (2017) Escape from the factory of the robot monsters: Agents of change. Team Performance Management: An International Journal, 23 (1/2), 96-108.

Richards, D. & Stedmon, A. (2016) To delegate or not to delegate: A review of control frameworks for autonomous cars. Applied Ergonomics, 53, 383-388.

Richards, D., Stedmon, A., Shaikh, S., & Davies, D. (2014) Responding to disaster using autonomous systems. The Ergonomist (Special Feature), No. 534, December 2014.

Simon, H. (1955) A behavioural model of rational choice. Quarterly Journal of Economics, 69, 99-118.

Weimer,C.W., Miller, J.O., & Hill, R.R. (2016) Agent-based modelling: An introduction and primer. Proceedings of the 2016 Winter Simulation Conference, INFORMS, December 11-14, USA: Washington D.C.

Wiedenbeck, S. (1999) The use of icons and labels in an end user application program: An empirical study of learning and retention. Behaviour & Information Technology,18(2), 68-82.

Wilson, M. (2002) Six views of embodied cognition. Psychonomic Bulletin & Review, 9(4), 625-636.

Zeng, Y., Jouandeau, N., & Cherif, A. A. (2013) A survey and analysis of multi-robot coordination. *International Journal of Advanced Robotic Systems*, vol. 10, 2013.