Shared Mental Models as a Way of Managing Transparency in Complex Human-Autonomy Teaming

Gabriele Scali and Robert D. Macredie
Computer Science Department, Brunel University, Kingston Lane, London, U.K.

Keywords: Human-Autonomy Teaming, Human-Agent Collaboration, Agent Transparency, Shared Mental Models, Dynamic Environments, Time Pressure, Environment Complexity.

Abstract: This paper argues that because of the cognitive and communication limitations of human and autonomous agents engaged in Human-Autonomy Teaming within dynamic environments, various external factors, which can be classified collectively as environment complexity, set boundaries to the effectiveness of strategies for agent transparency – that is, the ability of autonomous agents to make human actors aware of their goals, actions, reasoning, and expectations of future states. Understanding the mechanisms by which changes in environment complexity affect transparency, and the conditions in which it can be disrupted, can help researchers to better frame the results of existing and future studies on transparency and, in turn, inform the development of strategies to modify autonomous agents' behaviour to maintain transparency under different environment conditions. It is proposed that one such strategy could be the adjustment of the level of abstraction of the shared mental model adopted by the team as the common ground for communication so as to keep the amount of information that is exchanged manageable within human cognitive limitations.

1 INTRODUCTION

Improvements in capabilities of Artificial Intelligence (AI) create opportunities for autonomous agents to be deployed in increasingly diverse real-world work environments as partners in mixed human-agent teams (Sycara, 2002). These situations are the subject of interdisciplinary research into Human-Autonomy Teaming (HAT), which investigates the challenges related to human collaboration with autonomous agents towards the achievement of common objectives (Christoffersen and Woods, 2002; Hoc, 2000; How, 2016; Shively et al., 2018).

In order to be said to participate in a team as a true partner, an agent must be autonomous, which means it being able to generate its own goals and free to act on them (Luck and D'Inverno, 1995). To be autonomous, agents must be capable of surviving in their environment (be viable), they must not need help in performing their tasks (be self-sufficient), and they must set their own goals and make their own plans (be self-directed). The above characterisation can only be meaningful when referred to a specific context of activity (Bradshaw et al., 2013; Kaber, 2018). Throughout this paper, we refer to that context as the 'HAT environment', 'operational environment' or just 'environment'.

Agents involved in HAT are not bound by dependence relationships, as it is the case in supervisory control situations (Sheridan, 2012). Therefore, to be effective team mates, they must behave collaboratively (Bellamy, 2017; Klein et al., 2004). A key aspect of doing so is to remain transparent. This paper adopts the account of transparency proposed by the Situation Awareness-Based Agent Transparency (SAT) framework (Chen and Barnes, 2014), which defines transparency as the ability of an agent to make another aware of their goals, actions, reasoning, and expectations of future states. In order to do so, agents have to: select the information they intend to communicate; choose an appropriate time to communicate it; choose an appropriate channel to communicate it; decide when it is appropriate to repeat it; and decide when to communicate updates and confirmations.

Transparency is therefore to be understood as a quality of these actions and decisions, hingeing on humans and autonomous agents being able to share an understanding of the situation and of the mechanisms and rules governing it.

Existing research on transparency in HAT has focused on defining it as a construct and on

Scali, G. and Macredie, R.
Shared Mental Models as a Way of Managing Transparency in Complex Human-Autonomy Teaming.
DOI: 10.5220/0007484801530159
In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2019)*, pages 153-159
ISBN: 978-989-758-354-4

manipulating its level to study its effect on team performance (Chen et al., 2017; Stowers et al., 2017; Wohleber et al., 2017; Wright et al., 2016). Research is lacking, however, into the factors and mechanisms affecting the achievement of transparency itself, and consequently how to maintain it.

This paper proposes firstly that environment complexity affects agent transparency, as the demands it creates test the limits of agents' cognitive and communication abilities. This determines boundaries within which certain strategies for achieving transparency work effectively. When complexity exceeds these boundaries, it may be necessary for the agent to make adjustments in order to maintain transparency. Our second proposition is that one such adjustment can be made in regards to the level of abstraction of the Shared Mental Model (SMM) on which communication and understanding of the situation is grounded.

As an example let us consider a scenario in which a human pilot is teamed with a synthetic agent navigator, with the joint goal of visiting a certain number of waypoints that become known to the navigator over the length of the mission. The task of the pilot is to drive a vehicle and to negotiate the uncertainties of the terrain. The task of the navigator is to interpret incoming information and to relay it to the pilot so as to direct them to visit the waypoints in the most efficient way. With a low number of waypoints and a slow rate of arrival of new ones, the navigator communicates the exact position of each waypoint and the order in which it intends to visit them. The shared mental model is one of points on a map. If the number of locations and the rate of their arrival increase, at a certain point it would become difficult to maintain an effective communication between navigator and pilot: transparency would break down (our first hypothesis). The navigator agent may then choose to switch to a more abstract mental model, based on the density of locations to visit on the map, along with the specific position of only the waypoint to reach next. As long as both mental models have been practised in advance, and are equally familiar to the pilot, this may allow the team to re-establish transparency in the changed, more challenging, conditions (our second hypothesis).

The remainder of this paper will discuss complexity factors of dynamic environments and how they can hinder transparency, briefly introduce SMMs and their role in maintaining transparency and conclude by proposing a possible approach to mitigating this effect by adjusting the level of abstraction of the SMM.

2 COMPLEXITY IN DYNAMIC ENVIRONMENTS

An increasing number of useful applications of HAT are possible in dynamic environments (Bainbridge, 1997; Hoc, 1993; Russell and Norvig, 2009), characterised by the possibility for system changes to occur independently of an agent's actions, owing to spontaneously-occurring events or to actions by agents outside the team. The uncertainty about future states and action outcomes, together with the inherent variability of context, makes applications in dynamic environments the most challenging for HAT (Kaber, 2018).

For example, in a chemical processing plant, machines can break down or availability of certain resources may vary due to provisioning fluctuations; in a Command and Control (C2) application, an adversary may try to impede the operations, or visibility may change; and in an Unmanned Vehicle (UxV) scenario, interfering traffic, shifting weather conditions and mechanical problems may all occur independently of the vehicle's actions, and affect their precise outcome as well as the choice of best course of action.

Environment complexity can vary between or within its instantiations. For example, an Air Traffic Control (ATC) system may go through periods of low and high traffic (number of entities), as well as situations when flights are on schedule and there is no need to hurry, and others when there is a need to recuperate delays (time pressure). This determines that an autonomous agent operating in that environment will be faced with maintaining transparency under possibly very different conditions.

While environment complexity factors, characterised sometimes as task features and constraints of the operational environment, have been found in previous research to impact interaction with automation (Mosier et al., 2013), their effect on HAT transparency has not been examined. Several frameworks have, though, been proposed to categorise factors contributing to complexity of environments and of tasks performed within them (Ham et al., 2011; Liu and Li, 2011), an exhaustive review of which is beyond the scope of the present work. For this research, we focus on three of the most commonly-cited complexity factors and discuss how they can affect agent transparency when present in a HAT environment. In particular we consider: time pressure (Edland and Svenson, 1993; Liu et al., 2016); predictability

(Mosier et al., 2013); and number of entities and possible courses of action (Park and Jung, 2007).

3 HOW TIME PRESSURE, PREDICTABILITY AND NUMBER OF ENTITIES AFFECT TRANSPARENCY

That agents can be autonomous does not mean that they do not differ significantly from humans in attitudinal capabilities. The observation that humans generally have better soft skills and adaptability, while agents are able to account for, and process, more information quicker and are somewhat limited in their capacity of action in the physical world, can be traced back to the classic Fitt's list (Fitts et al., 1951). Although the list may require some adjustments owing to technological advancements since its inception, the basic observation remains valid that, for the moment, machines and humans have largely differing abilities. In particular, the ability of AI-based systems to process much more information than the human mind, along with their computational advantage, is likely to determine a divergence of intelligibility between humans and agents. As environment complexity increases it is not possible to expect that humans can be made aware of everything an agent perceives, does and reasons (Miller, 2014).

3.1 Time Pressure

One of the commonly-cited contributing factors of complexity is time pressure, which decreases the time available to provide and understand explanations. Examples of highly time-pressured scenarios include search and rescue, operating rooms, command and control, sport competitions, and many others.

In settings where time pressure is not a driving issue, agents have the option to slowly relay all of the necessary information, provide detailed explanations, suggest possible courses of action, and then take the backseat in decision making. Their decisions can be vetted, understood or questioned by human actors before any action is taken. This leads to scenarios of classic Human-Automation Interaction, with the agent losing its autonomy in decision-making and working instead as an advisor.

Where, however, the scenario is governed by time pressure, the dynamic of interaction changes: agents, with their superior computational speed and ability to handle many concerns at once, are able to cope with time pressure well beyond the point where human actors become helpless. There is, though, less time to exchange information and to understand the agent's decisions in depth; as such, issues of trust come to the fore. Time pressure thus generates a requirement for a higher throughput in exchanging and processing information about the agent's state, plans and predictions. Since the cognitive abilities of humans are fixed, the only ways to manage this are compression or omission of information.

3.2 Predictability

When a system is predictable by an agent but not by a human, there is an asymmetry of information, which in turn makes the agent's actions less intelligible. In addition, the communication of expectations not corresponding to the current perception of the human actor can generate surprise or exacerbate issues of trust. For example, while some of the events within dynamic environments can be fundamentally unpredictable, others are opaque to a human, but probabilistically approachable for Artificial Intelligence (AI), in particular with Machine Learning (ML). In other words, these environments provide the opportunity to 'shine' for agents that can, in real time, make better predictions or calculate according to better models than humans are able to. Generally, these are hard-to-explain mathematical models, however, and even more so in real time situations. Explainable AI (XAI) (Adadi and Berrada, 2018) is investigating ways for Artificial Intelligence to communicate the 'reasoning' behind its predictions and decisions by using explanation interfaces using techniques borrowed from research in recommender systems (Pu and Chen, 2006), but doing so is generally feasible only in offline situations, in which time is not a factor.

3.3 Number of Entities

The number of items and relationships to account for in an environment directly generate cognitive demand: systems can easily become so complex that their scale and intricacy prevent humans from fully understanding them; this is the case in any sufficiently advanced work of ingenuity, from skyscrapers to microprocessors, as well as in large socio-technical systems, like a hospital. The same can be said for the complexity of reasoning, many examples of which can be found in the current literature about XAI.

Accounting for more entities individually requires a larger mental model. While this is not generally a problem for agents, it rapidly becomes one for humans. Once the mental models diverge, communication breaks down, and it becomes hard for agents to describe their state and their actions in a way that the human actor will understand – creating a breakdown in transparency.

Having outlined how complexity factors of dynamic environments can contribute to the breakdown of HAT transparency, it is important to look at the concept of Shared Mental Models, since it is through them that human actors understand and communicate, or otherwise fail to, an agent's actions in such an environment.

4 SHARED MENTAL MODELS

Shared Mental Models (Scheutz et al., 2017) are knowledge structures that simplify reasoning about a certain system (all models are simplifications that maintain some properties and relationships while losing others, and they exist for certain practical purposes). In particular, a model is shared so that the parties using it can perform the same reasoning, simplifications and assumptions to communicate or collaborate (language itself is dominated by models).

The adoption of a SMM (Cannon-Bowers et al., 1993; Stubbs et al., 2007) and the careful choice of the content and timing of their communication (Bindewald et al., 2014; Goodman et al., 2016) are critical mechanisms for agent transparency, as they provide the anchoring for the information being communicated.

An important feature of mental models is their level of abstraction. Reduction and synthesis are two ways to make models more abstract. In turn, a more abstract description requires less data and is easier to summarise. As an example, it is possible to talk about Italy as being the shape a boot – a rather abstract model – yet, as necessary, one may refer instead to a map fitting the page of a book, or to an accurate digital map, describing every street in the country. Each is preferable to support different tasks and contexts of use. The first one when describing to a friend which part of the country one visited; the second to show the administrative regions; and the third to draw an itinerary from a hotel to a museum. The different models are not interchangeable, therefore it is critical to transparency that agents use a model with the appropriate level of abstraction to optimise mutual understanding within their current context.

Although most people have an intuitive sense of a model's level of abstraction, a few formalisations have been proposed (Hayakawa, 1949; Rasmussen, 1979; Sheridan, 2017; St-Cyr and Burns, 2001). Rasmussen's, in particular, neatly provides a powerful taxonomy: Models of Physical Form; Models of Physical Function; Models of Functional Structure; and Models of Abstract Function, that generalises well across domains.

In everyday interactions, people commonly use SMM at different levels of abstraction to refer to the same systems. For example, a car engineer may think of an engine in terms of thermodynamics and materials (physical form and function) when talking about his work, but in terms of elasticity, fun or power, when explaining the car to a friend (abstract function). The curricula in computing have recognised the ability of dealing with abstractions as one of the fundamental computational skills (Grover and Pea, 2013), and levels, or layers, of abstraction are a fundamental concept in computing.

5 ADJUSTMENT OF THE LEVEL OF ABSTRACTION OF THE SHARED MENTAL MODEL TO PRESERVE TRANSPARENCY

Given that transparency is seen as a prerequisite for HAT effectiveness (Chen et al., 2018; Christoffersen and Woods, 2002), it is desirable for an autonomous agent to be aware of the current level of complexity and to adapt its strategy to maintain it. While adaptive strategies are not new in agent computing – a rich tradition of research exists in regards to adjustable autonomy (Bradshaw et al., 2003; Johnson et al., 2011) – the focus of that research is on adjusting the Level of Automation (LOA) in semi-autonomous systems. We propose, instead, to investigate how an autonomous agent may adjust the level of abstraction of SMM to maintain transparency while remaining fully autonomous. As we have seen in the analysis of how complexity affects transparency, the main disruption happens in regards to more information having to be conveyed, and understood, or less time to do so. In other words the main limiting factor of transparency is throughput. As we have seen, abstraction of a model is a simplification by way of compression or reduction of information. As a result, we propose that when a dynamic environment becomes more complex, and the information to transmit becomes too much, thus breaking down transparency, it is possible to repair it by adjusting the level of abstraction of the SMM so that the amount of information that must be communicated remains manageable for a human actor.

To do so, agents must continuously assess the complexity of the environment by monitoring the factors contributing to it – for example by keeping a count of how many entities are present. Work in adjustable autonomy can inform the present research in regards to the strategies used by agents to detect and model the conditions that trigger the adjustments, this could include counting entities, keeping track of rates of change of the environment, and keeping track of the occurrence of unexpected events. When an agent decides that it must switch the level of SMM, it must do so in a way that is clear and does not cause loss of Shared Situation Awareness (SSA) (Grimm et al., 2018). One way of doing this would be to make sure, prior to their deployment, that human agents are equally familiar with the different mental models they will encounter, for example by use of training; other ways to prevent loss of SSA would be to mark the switch through explicit communication and to design the different levels to be clearly distinguishable. Another concern in turning to a more abstract mental model is that, by definition, some of its ability to carry detailed information is going to be lost, and therefore it becomes crucial to establish what trade-off is most advantageous between loss of transparency and loss of information.

6 SUMMARY AND FURTHER WORK

In this paper we have examined the relationship between the complexity of dynamic environments and transparency in HAT, and have argued that the effects of complexity occur mainly as a consequence of the demands that complexity puts on throughput of communication between autonomous agents and human actors, as well as on the ability of humans to process larger amounts of information. We have presented the concept of SMM and put forward that varying the level of abstraction of the SMM may mitigate the disruptive effects of complexity on HAT transparency, by allowing the re-establishment of a common ground in terms that can be understood and communicated effectively under the new complexity conditions. Finally, we have highlighted some ways in which this could negatively affect SSA.

Our current research programme is directed at testing the above hypotheses. To do so we intend to build a virtual environment for HAT in which factors of complexity and SMM can be manipulated, and to investigate ways of measuring transparency within it. The overall objective of the research is therefore to compare transparency in situations in which complexity is increased with the SMM unchanged, to other situations in which the SMM is adjusted to compensate the increase.

REFERENCES

Adadi, A. and Berrada, M. (2018). Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6. pp. 52138 - 52160.

Bainbridge, L. (1997). The change in concepts needed to account for human behavior in complex dynamic tasks. *IEEE Transactions on Systems, Man, and Cybernetics Part A:Systems and Humans*, 27(3). pp 351-359.

Bellamy, R. (2017). Human-Agent Collaboration: Can an Agent be a Partner? In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '17*. New York, New York, USA: ACM Press. pp. 1289–1294.

Bindewald, J., Miller, M. and Peterson, G. (2014). A function-to-task process model for adaptive automation system design. *International Journal of Human Computer Studies*, 72(12), pp. 822–834.

Bradshaw, J., Hoffman, R., Woods, D. and Johnson, M. (2013). The seven deadly myths of autonomous systems. *IEEE Intelligent Systems*, 28(3), pp. 54–61.

Bradshaw, J., Sierhuis, M., Acquisti, A., Feltovich, P., Hoffman, R., Jeffers, R. and Van Hoof, R. (2003). Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications. In: *H. Hexmoor, C. Castelfranchi and R. Falcone, eds., Agent Autonomy*. Boston, MA: Springer US. pp. 243–280.

Cannon-Bowers, J., Salas, E. and Converse, S. (1993). Shared Mental Models in Expert Team Decision-Making. In: *N. Castellan, ed., Individual and Group Decision Making: Current Issues*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Chen, J. and Barnes, M. (2014). Human-Agent teaming for multirobot control: A review of human factors issues. *IEEE Transactions on Human-Machine Systems*, 44(1), 13–29.

Chen, J., Barnes, M., Selkowitz, A. and Stowers, K. (2017). Effects of Agent Transparency on human-autonomy teaming effectiveness. In: *2016 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2016 - Conference Proceedings*. IEEE. pp. 1838–1843.

Chen, J., Lakhmani, S., Stowers, K., Selkowitz, A., Wright, J. and Barnes, M. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science*, 19(3), pp. 259–282.

Christoffersen, K. and Woods, D. (2002). How to Make Automated Systems Team Players. *Advances in Human Performance and Cognitive Engineering Research*, 2, pp. 1–13.

Edland, A. and Svenson, O. (1993). Judgment and Decision Making Under Time Pressure. In: *Time Pressure and Stress in Human Judgment and Decision Making*. Boston, MA: Springer US. pp. 27-40.

Fitts, P., Viteles, M., Barr, N., Brimhall, D., Finch, G., Gardner, E. and Stevens, S. (1951). Human Engineering for an Effective Air-Navigation and Traffic-Control System. Oxford, England: National Research Council, Div. of.

Goodman, T., Miller, M., Rusnock, C. and Bindewald, J. (2016). Timing within human-agent interaction and its effects on team performance and human behavior. In: *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support* (CogSIMA). IEEE. pp. 35–41.

Grimm, D., Demir, M., Gorman, J. and Cooke, N. (2018). Team Situation Awareness in Human-Autonomy Teaming: A Systems Level Approach. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1), pp. 149–149.

Grover, S. and Pea, R. (2013). Computational Thinking in K-12: A Review of the State of the Field. *Educational Researcher*, 42(1), pp. 38–43.

Ham, D., Park, J. and Jung, W. (2011). A framework-based approach to identifying and organizing the complexity factors of human-system interaction. *IEEE Systems Journal*, 5(2), pp. 213–222.

Hayakawa, S. (1949). Language in Thought and Action. Oxford, England: Harcourt, Brace.

Hoc, J. (1993). Some dimensions of a cognitive typology of process control situations. *Ergonomics*, 36(11), pp. 1445–1455.

Hoc, J. (2000). From human-machine interaction to human-machine cooperation. *Ergonomics*, 43(7), pp. 833–843.

How, J. (2016). Human-autonomy teaming. *IEEE Control Systems*, 36(2), pp. 3–4.

Johnson, M., Bradshaw, J., Feltovich, P., Jonker, C., Van Riemsdijk, M. and Sierhuis, M. (2011). The fundamental principle of coactive design: Interdependence must shape autonomy. *Lecture Notes in Computer Science, 6541 LNAI*. Springer, Berlin, Heidelberg. pp. 172–191.

Kaber, D. (2018). A conceptual framework of autonomous and automated agents. *Theoretical Issues in Ergonomics Science*, 19(4), pp. 406–430.

Klein, G., Woods, D., Bradshaw, J., Hoffman, R. and Feltovich, P. (2004). Ten challenges for making automation a team player in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), pp. 91-95.

Liu, D., Peterson, T., Vincenzi, D. and Doherty, S. (2016). Effect of time pressure and target uncertainty on human operator performance and workload for autonomous unmanned aerial system. *International Journal of Industrial Ergonomics*, 51, pp. 52–58.

Liu, P. and Li, Z. (2011). Toward understanding the relationship between task complexity and task performance. *Lecture Notes in Computer Science, 6775 LNCS*. pp. 192–200.

Luck, M. and D'Inverno, M. (1995). A Formal Framework for Agency and Autonomy. In: *First International Conference on Multiagent Systems*. AAAI Press MIT Press. pp. 1–7.

Miller, C. (2014). Delegation and transparency: Coordinating interactions so information exchange is no surprise. *Lecture Notes in Computer Science, 8525 LNCS*. Springer, Cham. pp. 191–202.

Mosier, K., Fischer, U., Morrow, D., Feigh, K., Durso, F., Sullivan, K. and Pop, V. (2013). Automation, task, and context features: Impacts on pilots' judgments of human-automation interaction. *Journal of Cognitive Engineering and Decision Making*, 7(4), pp. 377–399.

Park, J. and Jung, W. (2007). A study on the development of a task complexity measure for emergency operating procedures of nuclear power plants. *Reliability Engineering and System Safety*, 92(8), pp. 1102–1116.

Pu, P. and Chen, L. (2006). Trust Building with Explanation Interfaces. In: *Proceedings of the 11th international conference on Intelligent User Interfaces (IUI '06)*. ACM, New York, NY, USA, pp. 93-100.

Rasmussen, J. (1979). On the Structure of Knowledge - a Morphology of Mental Models in a Man- Machine System Context. Risö-M 2192, Risö National Laboratory, Roskilde, Denmark.

Russell, S. and Norvig, P. (2009). Artificial Intelligence: A Modern Approach, 3rd edition. Pearson.

Scheutz, M., DeLoach, S. and Adams, J. (2017). A Framework for Developing and Using Shared Mental Models in Human-Agent Teams. *Journal of Cognitive Engineering and Decision Making*, 11(3), pp. 203–224.

Sheridan, T. (2012). Human supervisory control. In *G. Salvendy, ed., Handbook of human factors and ergonomics 4th ed*. Hoboken, NJ, USA. John Wiley. pp. 990–1015.

Sheridan, T. (2017). Musings on Models and the Genius of Jens Rasmussen. *Applied Ergonomics*, 59, pp. 598–601.

Shively, R., Lachter, J., Brandt, S., Matessa, M., Battiste, V. and Johnson, W. (2018). Why human-autonomy teaming? In: *Advances in Intelligent Systems and Computing (Vol. 586)*. Springer Verlag. pp. 3–11.

St-Cyr, O. and Burns, C. (2001). Mental models and the abstraction hierarchy: Assessing ecological compatibility. In: *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 45(4), pp. 297–301.

Stowers, K., Kasdaglis, N., Rupp, M., Chen, J., Barber, D. and Barnes, M. (2017). Insights into human-agent teaming: Intelligent agent transparency and uncertainty. In: *Advances in Intelligent Systems and*

Computing (Vol. 499). Springer, Cham. pp. 149–160.

Stubbs, K., Wettergreen, D. and Hinds, P. (2007). Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22(2), pp. 42–50.

Sycara, K. (2002). Integrating Agents into Human Teams. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 46(1), pp. 413–417.

Wohleber, R., Stowers, K., Chen, J. and Barnes, M. (2017). Effects of agent transparency and communication framing on human-agent teaming. In: *2017 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2017 (Vol. 2017–January)*. pp. 3427–3432.

Wright, J., Chen, J., Barnes, M. and Hancock, P. (2016). Agent Reasoning Transparency's Effect on Operator Workload. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 60(1), pp. 249–253.