

Shared Mental Models as A Way of Managing Transparency in Complex Human-Autonomy Teaming: A Position Paper

Gabriele Scali¹ and Robert D. Macredie¹

¹*Computer Science Department, Brunel University, Kingston Lane, London, United Kingdom*

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 & 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 & 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, Hoffman, Woods, & Johnson, 2013; Kaber, 2017). 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, Woods, Bradshaw, Hoffman, & Feltovich, 2004). A key aspect of doing so is to remain transparent. This paper adopts the account of transparency proposed by Chen & Barnes (2014) in their Situation Awareness based Agent Transparency (SAT) framework, 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

manipulating its level to study its effect on team performance (Chen, Barnes, Selkowitz, & Stowers, 2017; Stowers et al., 2017; Wohleber, Stowers, Chen, & Barnes, 2017; Wright, Chen, Barnes, & Hancock, 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 & 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, 2017).

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, Park, & Jung, 2011; P. Liu & 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 & Svenson, 1993; D. Liu, Peterson, Vincenzi, &

Doherty, 2016); predictability (Mosier et al., 2013); and number of entities and possible courses of action (Park & 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 & 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 & 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, DeLoach, & Adams, 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 – if not made of – models).

The adoption of a SMM (Cannon-Bowers, Salas, & Converse, 1993; Stubbs, Wettergreen, & Hinds, 2007) and the careful choice of the content and timing of their communication (Bindewald, Miller, & Peterson, 2014; Goodman, Miller, Rusnock, & Bindewald, 2016) are critical mechanisms of 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.

For example, it is possible to think of Italy as being more or less shaped like a boot – a rather abstract model – or to refer to a small map fitting on one page of a book, or even to an accurate digital map, describing every street in the country. Each one of them requires more data to be described, and is harder to summarise. And each is to be preferred to communicate in different 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. One could not use these models interchangeably. Using a model with the appropriate level of abstraction to optimise mutual HAT understanding within the current context, is therefore critical to transparency.

Although most people have an intuitive sense of a model's level of abstraction, a few formalisations exist (Hayakawa, 1949; Rasmussen, 1979; Sheridan, 2017; St-Cyr & 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 & 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 & 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, Demir, Gormana, & Cooke, 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., & Berrada, M. (2018). Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 1–1.
- Bainbridge, L. (1997, May). 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*.
- Bellamy, R. K. E. (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* (pp. 1289–1294). New York, New York, USA: ACM Press.
- Bindewald, J. M., Miller, M. E., & Peterson, G. L. (2014). A function-to-task process model for adaptive automation system design. *International Journal of Human Computer Studies*, 72(12), 822–834.
- Bradshaw, J. M., Hoffman, R. R., Woods, D. D., & Johnson, M. (2013). The seven deadly myths of “autonomous systems.” *IEEE Intelligent Systems*, 28(3), 54–61.
- Bradshaw, J. M., Sierhuis, M., Acquisti, A., Feltovich, P. J., Hoffman, R. R., Jeffers, R., Van Hoof, R. (2003). Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications. In H. Hexmoor, C. Castelfranchi, & R. Falcone (Eds.), *Agent Autonomy* (pp. 243–280). Boston, MA: Springer US.
- Cannon-Bowers, J. A., Salas, E., & Converse, S. (1993). Shared Mental Models in Expert Team Decision-Making. In N. J. Castellan (Ed.), *Individual and Group Decision Making: Current Issues*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Chen, J. Y. C., & Barnes, M. J. (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. Y. C., Barnes, M. J., Selkowitz, A. R., & 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 (pp. 1838–1843). IEEE.
- Chen, J. Y. C., Lakhmani, S. G., Stowers, K., Selkowitz, A. R., Wright, J. L., & Barnes, M. J. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science*, 19(3), 259–282.
- Christoffersen, K., & Woods, D. D. (2002). How to Make Automated Systems Team Players. *Advances in Human Performance and Cognitive Engineering Research*, 2, 1–13.
- Edland, A., & Svenson, O. (1993). Judgment and Decision Making Under Time Pressure. In *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 27–40). Boston, MA: Springer US.
- Fitts, P. M., Viteles, M. S., Barr, N. L., Brimhall, D. R., Finch, G., Gardner, E., ... Stevens, S. S. (1951). *Human Engineering For An Effective Air-Navigation And Traffic-Control System*.
- Goodman, T. J., Miller, M. E., Rusnock, C. F., & Bindewald, J. M. (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) (pp. 35–41). IEEE.
- Grimm, D. A., Demir, M., Gormana, J. C., & Cooke, N. J. (2018). Team Situation Awareness in Human-Autonomy Teaming: A Systems Level Approach. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1), 149–149.
- Grover, S., & Pea, R. (2013). Computational Thinking in K-12: A Review of the State of the Field. *Educational Researcher*, 42(1), 38–43.
- Ham, D. H., Park, J., & Jung, W. (2011). A framework-based approach to identifying and organizing the complexity factors of human-system interaction. *IEEE Systems Journal*, 5(2), 213–222.
- Hayakawa, S. J. (1949). *Language in Thought and Action*. Oxford, England: Harcourt, Brace.
- Hoc, J.-M. (1993). Some dimensions of a cognitive typology of process control situations. *Ergonomics*, 36(11), 1445–1455.
- Hoc, J.-M. (2000). From human-machine interaction to human-machine cooperation. *Ergonomics*, 43(7), 833–843.
- How, J. P. (2016, April). Human-autonomy teaming. *IEEE Control Systems*, 36(2), 3–4.
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., & Sierhuis, M. (2011). The fundamental principle of coactive design: Interdependence must shape autonomy. In *Lecture Notes in Computer Science* (Vol. 6541 LNAI, pp. 172–191). Springer, Berlin, Heidelberg.
- Kaber, D. B. (2017, November 27). A conceptual framework of autonomous and automated agents. *Theoretical Issues in Ergonomics Science*, pp. 1–25.
- Klein, G. A., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004, November). Ten challenges for making automation a “team player” in joint human-agent activity. *IEEE Intelligent Systems*.
- Liu, D., Peterson, T., Vincenzi, D. A., & 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, 52–58.
- Liu, P., & Li, Z. (2011). Toward understanding the relationship between task complexity and task performance. *Lecture Notes in Computer Science*, 6775 LNCS, 192–200.
- Luck, M. M., & D’Inverno, M. (1995). A Formal Framework for Agency and Autonomy. In *First International Conference on Multiagent Systems* (pp. 1–7). AAAI Press MIT Press.
- Miller, C. A. (2014). Delegation and transparency: Coordinating interactions so information exchange is no surprise. In *Lecture Notes in Computer Science* (Vol. 8525 LNCS, pp. 191–202). Springer, Cham.
- Mosier, K. L., Fischer, U., Morrow, D., Feigh, K. M., Durso, F. T., Sullivan, K., & 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), 377–399.
- Park, J., & Jung, W. (2007). A study on the development of a task complexity measure for emergency operating procedures of nuclear power plants. *Reliability Engineering & System Safety*, 92(8), 1102–1116.
- Pu, P., & 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, 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., & Norvig, P. (2009). Artificial Intelligence: A Modern Approach, 3rd edition. Pearson.
- Scheutz, M., DeLoach, S. A., & Adams, J. A. (2017). A Framework for Developing and Using Shared Mental Models in Human-Agent Teams. *Journal of Cognitive Engineering and Decision Making*, 11(3)
- Sheridan, T. B. (2012). Human supervisory control. In G. Salvendy (Ed.), *Handbook of human factors and ergonomics* (4th ed., pp. 990–1015). Hoboken, NJ, USA: John Wiley.
- Sheridan, T. B. (2017). Musings on Models and the Genius of Jens Rasmussen. *Applied Ergonomics*, 59, 598–601.
- Shively, R. J., Lachter, J., Brandt, S. L., Matessa, M., Battiste, V., & Johnson, W. W. (2018). Why human-autonomy teaming? In *Advances in Intelligent Systems and Computing* (Vol. 586, pp. 3–11). Springer Verlag.
- St-Cyr, O., & Burns, C. M. (2001). Mental models and the abstraction hierarchy: Assessing ecological compatibility. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 45(4), 297–301.
- Stowers, K., Kasdaglis, N., Rupp, M. A., Chen, J. Y. C., Barber, D., & Barnes, M. J. (2017). Insights into human-agent teaming: Intelligent agent transparency and uncertainty. In *Advances in Intelligent Systems and Computing* (Vol. 499, pp. 149–160). Springer, Cham.
- Stubbs, K., Wettergreen, D., & Hinds, P. J. (2007). Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22(2), 42–50.
- Sycara, K. (2002). Integrating Agents into Human Teams. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 46(1), 413–417.
- Wohleber, R. W., Stowers, K., Chen, J. Y. C., & Barnes, M. J. (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–Janua, pp. 3427–3432). IEEE.
- Wright, J. L., Chen, J. Y. C., Barnes, M. J., & Hancock, P. A. (2016). Agent Reasoning Transparency's Effect on Operator Workload. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 60(1), 249–253.