Multimodal, Adaptable, and Dynamic Human Autonomy Team Relationships

Daniel Barber, Lauren Reinerman-Jones, Ryan Wohleber
Soar Technology, Inc.
Orlando, FL
daniel.barber@soartech.com, lauren.reinerman@soartech.com, ryan.wohleber@soartech.com

Joseph B. Lyons
U.S. Air Force Research Laboratory
Wright-Patterson Air Force Base, OH
joseph.lyons.6@us.af.mil

ABSTRACT

The military operates in dynamic and complex task environments with inherent uncertainty that can make it difficult for human operators to form appropriate strategies when using autonomy. Uncertainty about system capability often results in disuse or misuse of autonomy, and as a result, performance and situation awareness may be compromised. A key component in successful teaming is trust. Methods of transparency are emerging to support human-autonomy teams (HATs) by communicating information about autonomy's actions, decisions, behaviors, and intentions to develop shared awareness and shared intent. Most efforts investigate transparency during task execution, but there are limits to how efficiently transparency content is communicated in these settings. Given that and the ability to process at high speed, HAT operation might prohibit transparency information delivery and operator perception and comprehension before actions must be taken. Consideration of transparency information during task performance can siphon valuable mental resources during high task load, and lack of available mental capacity during these periods limits the ability of operators to learn and retain insights about how the autonomy functions. Addressing operator limits, implementation of transparency beyond mission execution must be sought. This paper presents guidelines for implementing adaptable multimodal transparency across a HAT lifecycle (e.g., pre, during, and post task) through the integration of multiple methods of transparency using an illustrative example from the authors' current research within the counter small, unmanned air systems domain.

ABOUT THE AUTHORS

Dr. Daniel Barber is a Lead Scientist at SoarTech with over 15 years of experience conducting and applying interdisciplinary research within the fields of modeling and simulation, augmented cognition, physiological assessment, robotics, machine learning, human-agent collaboration, control systems, path-planning, computer vision, communication frameworks, and environment modeling.

Dr. Reinerman-Jones is a Senior Scientist at SoarTech with over 20 years of experience leading multidisciplinary teams in science and technology funded efforts focused on assessment for explaining, predicting, and improving human performance and system design. She has published and presented in 9 countries on topics in autonomy, human-robot teaming, and AI.

Dr. Wohleber has over 12 years' experience in theoretical and applied human factors research investigating cognitive, emotional, and motivational factors in task performance. His work is documented in over 40 publications and has informed training, task design, augmented cognition systems, AI-powered UIs, cyber security practices, and other applications.

Dr. Lyons is a Principal Research Psychologist with over 16 years experience working with the Air Force Research Laboratory (AFRL). His work has focused on trust in autonomy, human-autonomy teaming, and human factors. He is an AFRL Fellow and a Fellow of the American Psychological Association and Society for Military Psychologists.

Multimodal, Adaptable, and Dynamic Human Autonomy Team Relationships

Daniel Barber, Lauren Reinerman-Jones, Ryan Wohleber
Soar Technology, Inc.
Orlando, FL
daniel.barber@soartech.com, lauren.reinerman@soartech.com, ryan.wohleber@soartech.com

Joseph B. Lyons
U.S. Air Force Research Laboratory
Wright-Patterson Air Force Base, OH
joseph.lyons.6@us.af.mil

INTRODUCTION

The military operates in dynamic and complex task environments with inherent uncertainty that can make it difficult for human operators to form appropriate strategies when using autonomy. Uncertainty about system capability often results in disuse or misuse of autonomy (Lee & See, 2004), and as a result, performance and situation awareness can be compromised. A key component in how successfully humans interact with autonomy is trust (Sheridan, 1988; Jiun-Yin Jian et al., 1998; Parasuraman & Wickens, 2008). Methods of transparency have emerged to support building trust in human autonomy teams (HATs) by communicating information about autonomy's actions, decisions, behaviors, and intentions to develop shared awareness and shared intent. Most efforts to investigate approaches to implementation of this communication have focused on providing transparency during task execution.

There are limits to how efficiently transparency can be communicated during task execution (Miller, 2021). Artificial intelligence systems can parse complex situations at a rapid speed, which in a HAT operation might prohibit comprehensive information delivery and operator perception and comprehension of the information about system behaviors, states, plans, and rationales before action must be taken. Recent DARPA exercises with swarms of drones exemplified this inability of a human operator to respond appropriately at the timescales of advanced artificial intelligence systems (Knight, 2021). Moreover, consideration of transparency information during task performance can siphon valuable mental resources during high task load, (Miller, 2021) and lack of spare mental capacity during these periods limits the ability of operators to learn and retain insights about how the autonomy functions. Partnering humans with autonomy enables force multiplication and superhuman response times to counter and defeat new and emerging threats, but to do this, we must look across the lifecycle of HAT for additional opportunities outside of mission execution to implement transparency methods to address challenges to information delivery and comprehension. Two important questions that a transparency capability should address are:

1. What is the impact of applying transparency methods across the HAT lifecycle to common ground, trust, communication, and reliance (e.g., pre, during, and post task)?
2. How do HAT interactions shape and change transparency communication needs over time?

As cumulative experiences form and adjust human mental models and schemas of autonomy, information needs during task execution may become more subtle and communication more routine. Thus, at the appropriate time in the HAT lifecycle, only strategic injections of transparency for critical clarifications may be necessary, allowing displayed information to be minimized and the task space decluttered (contrast to transparency visual representations in current approaches). The next critical step in developing a coherent and comprehensive transparency strategy is to integrate multiple methods that structure transparency before, during, and after task execution. Only by integrating theory across this lifecycle of transparency can a complete strategy for the development, maintenance, and repair of mental models and schemas fostering effective teaming relationships be realized.

This paper presents ongoing work in the development of Multimodal, Adaptable, and Dynamic Human Autonomy Team Relationships (MAD-HATR) that will implement adaptable multimodal transparency throughout a HAT lifecycle. The approach taken under this effort applies models of transparency to an existing autonomy capability matched throughout a HAT lifecycle to improve common ground, calibrate trust, maximize communication efficiency, and optimize autonomy use and reliance. The HAT lifecycle consists of enhanced training followed by repeated pre (e.g., pre-mission brief), during, and post task (e.g., debrief, after action review) phases of interaction. How models of transparency are integrated and matched to the HAT lifecycle is described, in addition to an experimental design to assess benefits of including transparency across HAT lifecycle phases.

BACKGROUND

The scientific literature currently identifies affordances for transparency before, during, and after task performance (Lyons, 2013; Miller, 2021), approaches for implementing transparency during task execution (Lyons, 2013; Chen et al., 2018; Vered et al., 2020), the mechanisms by which transparency affects HAT performance (transparency as a trust antecedent; Kaplan et al., 2021), transparency regarding human team member status (Lyons, 2013), function allocation (Bruni et al., 2007; Lyons, 2013) and more. While these theoretical underpinnings are robust, researchers and developers have yet to combine them into a single, comprehensive transparency communication strategy. A review of current practice reveals three additional gaps in theory-based transparency research and development. First, efforts to communicate transparency information have focused primarily on designing interfaces for use during task/mission execution, missing opportunities to establish common ground and correct mental models of autonomy across the HAT lifecycle described by some frameworks (Lyons, 2013). Second, implementation of transparency into autonomy interfaces has concentrated on a static, or fixed level of, visual representations (Bhaskara et al., 2020; Lakhmani et al., 2016; Mercado et al., 2016; Miller, 2021). Using a single perceptual channel (e.g., visual) limits the amount of information that can be communicated for both task work and teaming coordination, and a lack of redundancy increases risk of missed signals (Chhan et al., 2020). Multimodal communication across more than one sensory channel would allow for increased robustness and efficiency of interactions using redundant and non-redundant signals (Bischoff & Graefe, 2002; Burke et al., 2006). Though multimodal approaches to transparency design have been prescribed (Rafael, 2020), implemented in live systems (Hastie et al., 2018), and in one case undergone evaluation to improve transparency signaling (Chhan et al., 2020), to our knowledge no study has yet systematically evaluated multimodal approaches to theory-based transparency communication (c.f., Theodorou et al., 2017, p. 6). Third, the ability to flexibly select how one interacts with autonomy has demonstrable benefits in terms of time on task, task completion, and mission effectiveness (Taylor et al., 2015). While the use of adaptable autonomous systems is well studied (e.g., Calhoun et al., 2013), less well understood are the tradeoffs when users govern the *transparency* of autonomy teammates. Investigation of these tradeoffs are critical, as human control of transparency communication may be susceptible to and exacerbate trust miscalibration (Parasuraman & Riley, 1997).

APPROACH

The goal for MAD-HATR is to integrate multiple methods of transparency and operationalize them across a HAT lifecycle to measure any added benefits to operators. To meet this goal, there are three challenges that must be addressed which include: 1) how to integrate methods of transparency from the literature, 2) approaches to operationalize the transparency at different lifecycle phases, and 3) measurement of any benefits to the approach for the operator. This paper describes an approach to address these challenges using previously developed human-machine interface (HMI) and autonomy within the domain of counter small-unmanned air systems for the Air Force. It is expected that advancements made to enhance this testbed will support operators in developing robust mental models of autonomy when applied across the HAT lifecycle. The benefits of these mental models could include improved common ground, calibrated trust, maximized communication efficiency, and optimal autonomy reliance.

HAT Lifecycle

Our proposed HAT lifecycle begins with training, where the application of transparency can teach users more than basic "buttonology" and how to interact with autonomy, but more importantly the design, purpose, and intent of autonomy. After training, the lifecycle iterates and loops over the remaining phases beginning with a pre-mission brief to establish common ground, followed by task execution, and concludes with a post task such as an interactive debrief or after action review. The purpose for each phase is to support the user's development and maintenance of mental model(s) of the system/autonomy. Adaptable transparency is enabled with the user able to "increase" or "decrease" the amount of transparency information displayed at each lifecycle phase.

Over the course of the HAT lifecycle users experience multiple transparency interactions that facilitate their ability to correct mental models, anticipate future autonomy behaviors, reduce information needs during task execution, and calibrate trust resulting in a relationship where they appropriately rely on autonomy, Figure 1.

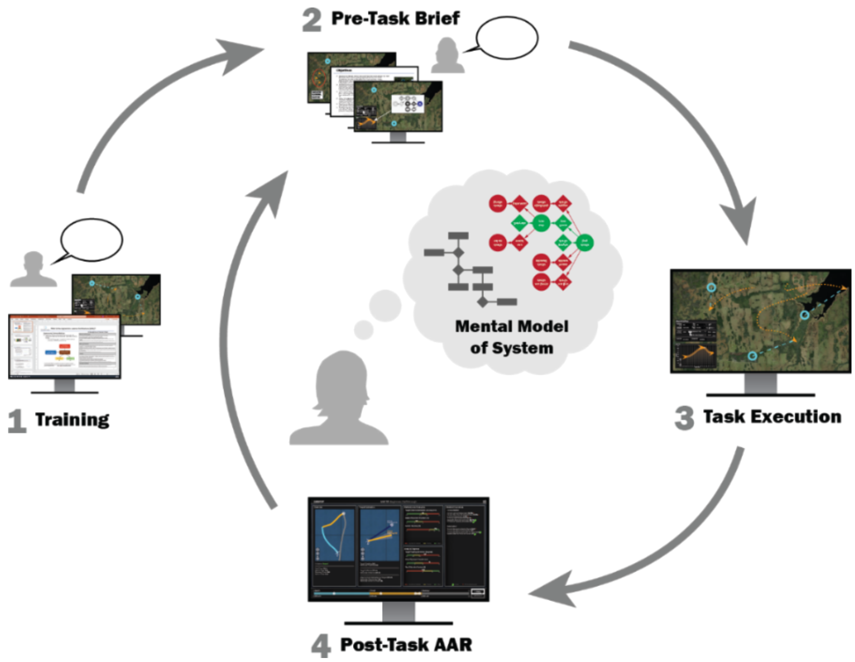

Figure 1. HAT lifecycle for autonomy users. During training 1) users learn about the design and underlying functions of the system to develop initial mental models of autonomy. After training, the lifecycle iterates and loops over three phases from pre-mission brief to establish common ground 2), during task execution 3), and post task (e.g., debrief, after action review). 4) During these phases human team members build common ground and calibrate trust over time through development of shared mental models.

Integrating Transparency Models

The first challenge to address is determining how to align different methods of transparency and appropriately match and apply them to the HAT lifecycle. The two most referenced models of transparency found in literature are Lyons' models of transparency for human-robot interaction and Chen et. al.'s situation awareness-based agent transparency (SAT) model, (Lyons, 2013; Chen et al., 2014). Lyons describes several factors important to effective human-robot teaming falling into two aspects: 1) the robot communicating information about its knowledge and views (robot-to-human), and 2) robot communication of awareness about human operator state (robot-of-human) (Lyons, 2013; Bhaskara et al., 2020). As an example for robot-to-human communication, Lyon's intentional model helps the user understand the design, purpose, and intent of a robot while the task model information can develop understanding of the task and how it is executed by the robot. At a more granular level, transparency categorized under Lyon's analytic model supports human understanding of the underlying principles that guide robot functions. Finally, the environment model provides users with awareness of task performance context and an understanding of the implications of that context.

The SAT model provides guidance on what information an agent should communicate to humans to promote transparency about the agent. Inspired by Endsley's (1995) theory of situation awareness (SA) and Lee and See's (2004) theory of human trust in automation, information is categorized into three levels. Level 1 regards what the agent is trying to achieve, Level 2 indicates why the agent is doing what it does, and Level 3 provides projections about what the robot expects should happen under current and hypothetical circumstances (Chen et al., 2014). The categorization of transparency content under the SAT model also supports implementation of adaptable transparency. Operators can therefore increase transparency to support their needs by increasing the transparency level of SAT content. The approach for this effort leverages multiple transparency methods to support the development and correction of each of Lyon's models of robot-to-human transparency (e.g., task model, analytical model) at different phases of the HAT lifecycle by incorporating transparency content organized and classified using the SAT model (Figure 2).

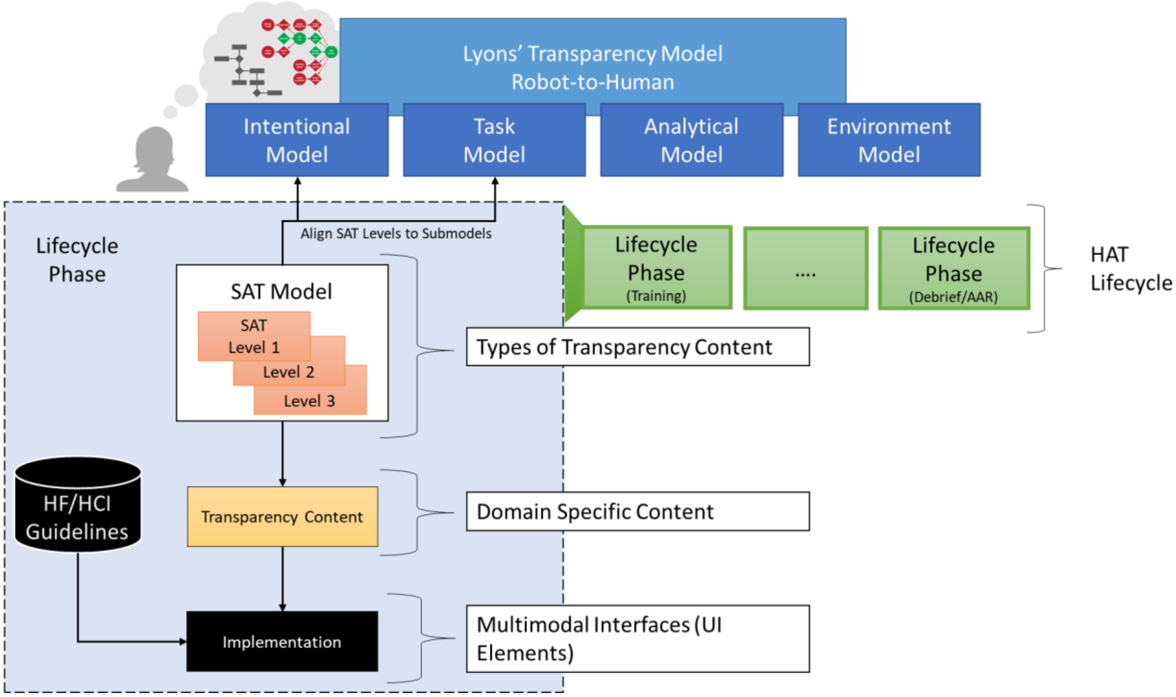

Figure 2. Approach to align SAT model with Lyon's robot-to-human transparency model across different phases of a HAT lifecycle.

As previously described, the HAT lifecycle includes four phases, with each phase providing opportunities to develop common ground, calibrate trust, reduce information needs, and improve reliance. For example, the training phase can facilitate development of users' intentional model of autonomy by providing insights about the system's design, purpose, and intent. Pre-mission briefs may support an operator's acquisition of content about how the autonomy approaches specific task performance (task models), what algorithms underpin task performance (analytical models), and the implications of these capabilities for the current tasking environment (environment model). This information supports user inference during task-execution pertaining to what the autonomy is currently doing and what autonomy may do in the future, reducing the need for clarification performing a mission. Finally, the after action review will review HAT performance, providing details about autonomy behaviors to support user review of autonomy at different points in the mission. These reviews will provide timely, situated insights to further support the development of Lyons' robot-to-human transparency models in human operators. Moreover, these after action reviews support repair or "synchronization" of mental models to recalibrate trust and, by extension, reliance on autonomy over time (Miller, 2021).

OPERATIONALIZING TRANSPARENCY

As illustrated in Figure 2, transparency content can be classified and organized based on the type of situation awareness and robot-to-human model(s) supported for each phase in the HAT lifecycle. However, to generate and deliver transparency information according to the guidelines of this integrated framework requires that these principles be integrated into components of an existing system (challenge two), which may also require the development of new algorithms. The transparency communication strategy is operationalized within the counter small unmanned air systems (sUAS) domain for this program. Under a previous effort for the Air Force, autonomy and a human machine-interface was developed for a system with the ability to propose courses-of-action to mitigate small unmanned air system threats using combinations of different effectors. This system was built upon the Air Force's Fusion software used in the development of the Intelligent Multi-Unmanned Vehicle Planner and Adaptive Collaborative/Control Technologies testbed, (Draper et al., 2017). In the current Counter-sUAS autonomy stack, rationale behind course of action options is based on rules of engagement and pairing of available resources to threats. Final ranking of courses of action is performed using weighted scoring of different factors (e.g., threat level, time). Although this provides a basic level of transparency content to an operator (e.g., allowed actions, ranked courses of action), there are gaps concerning what other data is needed and how to generate transparency content at multiple SAT levels (e.g., current state, rationale, predictions) to facilitate the construction of appropriate mental models over the span of HAT lifecycle interactions. For example, the current system does not provide the operator rationale for why a course of action is unavailable for a given sUAS. Moreover, under certain conditions, the autonomy may have the authority to enact a course of action to mitigate a threat on its own. Currently, this automatic action can confuse operators and impact their trust and reliance when they lack understanding of the conditions under which automatic actions may occur. Further, the Counter-sUAS platform provides many opportunities to apply transparency multimodally (e.g., visual display elements, speech dialogue) and can serve as the foundation for the generation of content to support transparency communication across each lifecycle phase.

To implement our theory-based transparency communication guidelines across the HAT lifecycle, current and potential capabilities of the Counter-sUAS system relevant to autonomy task performance and teaming were identified and the communications these capabilities afforded were categorized within each SAT level. For example, one potential capability identified to support a more theoretically complete transparency communication strategy was the inclusion of projections of future sUAS positions (affording SAT Level 3). Next, these capabilities and their prospective communications were aligned to Lyon's robot-to-human models (e.g., analytical model), with explanations as to how data from each capability can directly support each model. The third step was to identify how to instantiate prospective transparency content from these capabilities across each phase of the HAT lifecycle. Potential solutions included multiple methods of representing transparency content, including text, graphics, iconography, and speech, Table 1. It is important to note that how the system conducts a phase of the lifecycle will directly impact how transparency can be delivered. For example, if a pre-mission brief is provided as a five-part operations order, then the format will be primarily text. If that same brief is done using a visual display or other multimedia, then visual representations could be made available. Therefore, system designers must consider not just the SAT categorization of affordances and supported mental models, but the ability to support interaction modalities for each lifecycle phase, Figure 3.

Table 1. Example table of affordance data mapping Lyon's robot-to-human analytic model data to Chen's SAT Levels. Table categorizes the type of transparency content supported/available at SAT levels and recommendations/options for implementation. This example is for authorized and prohibited actions.

Transparency Content			
Affordance	**SAT 1**	**SAT 2**	**SAT 3**
Authorized Actions	List of currently authorized and prohibited actions for a sUAS.	Which rules-of-engagement determined the authorization or prohibition of actions for a sUAS.	Upcoming/future authorizations and prohibition based on a projection of future track state (e.g., position).
Transparency Implementation Options			
Affordance	**SAT 1**	**SAT 2**	**SAT 3**
Authorized Actions	1) Visual display panel listing authorized actions for each sUAS. 2) Ability to select sUAS on map using touch or mouse interaction to pull up a list or iconography indicating available actions. 3) speech interface to query system for authorized or prohibited actions for sUAS.	1) Additional information in visual display indicating the rules-of-engagement or conditional information pertaining to rules authorizing or prohibiting actions. 2) Additional iconography on visual display panel or when an operator selects a sUAS on a map display indicating relevant rules-of-engagement and factors for each authorized action. 3) Speech interface to query system for explanation for available action(s) or prohibition(s).	1) Visual display panel indicating future list of authorized actions for each sUAS at one or more time increments. 2) Ability to select sUAS on map using touch or mouse interaction to pull up a list or iconography indicating available future actions. 3) speech interface to query system for predicted authorized or prohibited actions for sUAS at a given future time.

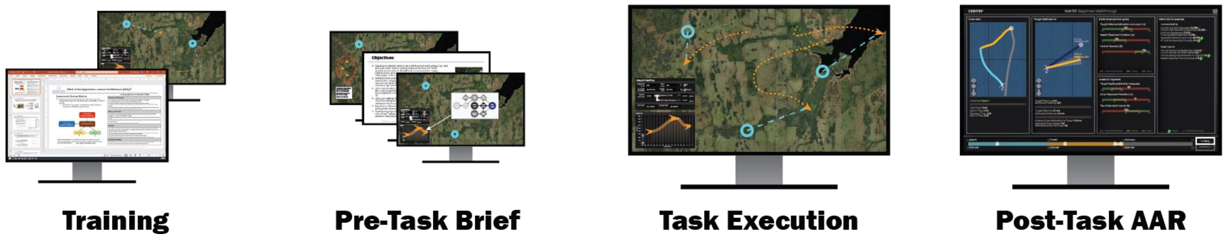

Figure 3. Examples of different interaction modalities and interfaces for HAT lifecycle phases, including Training, Pre-Task Brief, Task Execution (during mission), and Post-Task Debrief/After-Action Review (AAR).

ASSESSING THE BENEFIT OF TRANSPARENCY ACROSS HAT LIFECYCLE

The third challenge identified for this effort is to assess any benefit of transparency applied across the HAT lifecycle. Adding transparency to an existing or new system will likely result in additional funding requirements and time to deliver these modifications, and as such it is important to understand the benefits associated with these costs. Although prior work in the literature has shown the benefits of transparency during task execution, efforts have not evaluated any effects resulting from transparency communications in other phases of a HAT lifecycle or trends resulting from these communications in multiple iterations of one phase. Knowing these benefits will directly support requirements for future HAT systems, allowing designers to focus on transparency where it will have the most effect. To address this, a between-subjects manipulation of transparency inclusion on different phases of the HAT lifecycle over multiple iterations is proposed which can measure impact on common ground, trust, and reliance (mixed-model), Table 2.

Table 2. Proposed experimental design to evaluate transparency across a HAT lifecycle. Design is between-subjects manipulation of transparency inclusion on different phases with transparency included (T) or not included (NT).

	Lifecycle Phase			
HAT Lifecycle	**Training**	**Pre**	**During**	**Post**
1	NT	T	T	T
2	T	NT	T	T
3	T	T	NT	T
4	T	T	T	NT
5	T	T	T	T
6	NT	NT	NT	NT

The above experimental design ensures appropriate comparisons between groups in that it separates the impact of a lifecycle phase from the inclusion (or not) of transparency for that phase. Like Figure 1, the Training phase would occur one time, and then the remaining phases would loop over multiple iterations. Repeated iteration of the final three phases enables measurement of effects on common ground, trust, and reliance over time.

An added challenge for this design is identifying how to appropriately include or exclude transparency and how that manipulation aligns to robot-to-human transparency submodels. As an example, the task model supports user sensemaking of the actions of the autonomy within the context of its performance. The task model therefore provides information about the task being performed as perceived by the system, intelligence about the goals serviced by performing the task, system fitness for the task, and task progress. This information can help users evaluate whether the system holds the same mental model for the task and its implications and evaluate performance informed by the systems own tracking of the task. SAT data has long been used to convey information about the system relative to a task being performed by the system during performance. Examples include communicating mode, perception of task space (mission events that affect approach to tasking), fitness for task (plan recommendations and fitness of plan), task status progress, and system status (e.g., fuel level). If transparency support for the task model were to be manipulated as proposed in the experimental design, then removing transparency would mean removing key pieces of information an operator would need to perform a task. As a result, task model information can be increased to enhance transparency, but a base level is embedded and cannot be removed without degrading the overall task. Similarly, SAT level transparency could potentially be added to support development of other submodels (e.g., intentional, environmental) in some phases of the lifecycle, but not others.

Analytic Model

While task model frames functions and behaviors within a task performance context, prospectively, in real-time, and retrospectively, the analytic model of transparency is about communicating with the human counterpart understandable and actionable explanations and insights concerning the underlying processes that power the system. The analytic model can be considered the inner most layer covering the principles and even the specific algorithms that underpin system function. Whereas the task model requires heuristics and considerable guidance from trust attitudes to inform reliance, the analytic model is much more knowledge-based and can be scrutinized if necessary to come to a more objective conclusion. The analytic model can also inform expert-level decisions in exceptional situations but has only been a part of the transparency solution in practice, or completely ignored. Much of the existing research about the SAT model focuses on the task model level and has shown transparency information on that level to be very useful in calibrating reliance. However, there are instances when task level information is insufficient for understanding behavior, especially when that behavior is anomalous and difficult to derive from intended system functions, the systems behaviors or approaches to tasks, and performance on those tasks. These instances of failure of intuitive understanding of a system can have serious impacts on trust in the system and undermine reliance calibration. SAT information in support of the analytic model can focus on where more abstract levels of transparency fail to inform proper reliance and in the interest of minimizing mental load. Moreover, information supporting development of the analytic model, and ideally an intuitive understanding of autonomy, can best support a transparency inclusion manipulation as proposed where other models fail.

Analytic Model Relationship to SAT Model

For SAT Level 1, analytic model information refers to objective information about how a decision, behavior, function, etc. is calculated or derived (i.e., principles, algorithms). It might communicate what different modes or calculation methods are in play, but explanation of these modes or algorithms beyond their objective purpose and function would be considered SAT Level 2 information. SAT Level 1 information alone depends on the human to determine the implications of the underlying function and mechanisms as explained. Within the domain of Counter-sUAS, specific examples of SAT Level 1 information (e.g., status of a particular parameter, such as visibility) may allow human operators to make inferences based on their analytic model of the system, specifically how these parameters directly feed into autonomy algorithms and impact system outputs. For example, base defense zones defined on a map have characteristics associated with them that drive what actions may be authorized within them and therefore how algorithms underpinning automated effector behaviors might respond. Violations of these expectations will inform updates to the operator's analytic model of the system. SAT Level 1 may also involve direct communication of the algorithm, though in the artificial intelligence space, it is usually functional principles that inform the analytic model best, as calculation-level information is usually too complex or uninterpretable (e.g., deep learning algorithms).

Analytic model SAT level 2 information provides the logic or reasoning behind why the calculations, functional principles, or underlying mechanisms are employed – i.e., the reasoning. This might mean that depending on conditions, different algorithms are applied to a problem and SAT 2 level information is needed to provide the reasoning behind this shift in system strategy (SAT 1 level information would convey the shift itself). For example, a specific rule of engagement authorized or prohibited, and action is level 2, while the state of the action (authorized/prohibited) is level 1.

SAT level 3 information in support of the analytic model would be a matter of deconstructing projections to their component parts, and/or involve tweaking of models or offering alternative models based on analytic model levels. This can improve understanding of system function by demonstrating the expected consequences of an underlying mechanism. For example, projections might provide insights about different modes or simply illustrate an unexpected consequence of the system's programming. History of performance can illustrate this as well.

CONCLUSION

This paper describes the need for a theory-based approach to transparency communication across all phases of the HAT lifecycle (training, pre, during, and post-task) and guidelines for filling this gap. Specifically, the authors advocate for an approach that systematically integrates two prominent theories of transparency, Lyon's robot-to-human transparency model and Chen and colleagues' SAT model, to optimize communication strategies across lifecycle phases. This approach would leverage insights from over a decade of task-based transparency research that has utilized these models. To illustrate these proposed guidelines, we highlight ongoing work to integrate these models to inform the construction of a multimodal testbed to support transparency research in the counter-UAS domain and demonstrate how this integrated theoretical approach informs capability requirements. Finally, we prescribe an experimental design to address how to assess the impact of implementing theory-based transparency across iterations of the HAT lifecycle to support development of analytical models of autonomy teammates in human operators.

ACKNOWLEDGEMENTS

This work is supported, in part by the US Air Force Research Laboratory under contract number FA8650-22-C-6417. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the AFRL or the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

Bhaskara, A., Skinner, M., & Loft, S. (2020). Agent Transparency: A Review of Current Theory and Evidence. *IEEE Transactions on Human-Machine Systems*, *50*(3), 215–224. https://doi.org/10.1109/THMS.2020.2965529

Bischoff, R., & Graefe, V. (2002). Dependable multimodal communication and interaction with robotic assistants. *Proceedings. 11th IEEE International Workshop on Robot and Human Interactive Communication*, 300–305. https://doi.org/10.1109/ROMAN.2002.1045639

Bruni, S., Marquez, J. J., Brzezinski, A., Nehme, C., & Boussemart, Y. (2007). Introducing a human-automation collaboration taxonomy (hact) in command and control decision-support systems. *12th International Command and Control Research and Technology Symposium*. http://dodccrp.org/events/12th_ICCRTS/CD/html/papers/064.pdf

Burke, J. L., Prewett, M. S., Gray, A. A., Yang, L., Stilson, F. R. B., Coovert, M. D., Elliot, L. R., & Redden, E. (2006). Comparing the effects of visual-auditory and visual-tactile feedback on user performance: A meta-analysis. *Proceedings of the 8th International Conference on Multimodal Interfaces - ICMI '06*, 108. https://doi.org/10.1145/1180995.1181017

Calhoun, G. L., Draper, M. H., Miller, C., Ruff, H. A., Breeden, C., & Hamell, J. (2013). Adaptable automation interface for multi-unmanned aerial systems control: Preliminary usability evaluation. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *57*(1), 26–30. https://doi.org/10.1177/1541931213571008

Chen, J. Y. C., Lakhmani, S. G., Stowers, K., Selkowitz, A. R., Wright, J. L., & Barnes, M. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science*, *19*(3), 259–282. https://doi.org/10.1080/1463922X.2017.1315750

Chen, J. Y. C., Procci, K., Boyce, M., Wright, J., Garcia, A., & Barnes, M. (2014). Situation Awareness-Based Agent Transparency. *DTIC*. BASE. https://search.ebscohost.com/login.aspx?direct=true&db=edsbas&AN=edsbas.36C22AC4&site=eds-live&scope=site&custid=current&groupid=main&authtype=shib

Chhan, D., Scharine, A., & Perelman, B. S. (2020). *Human-Autonomy Teaming Essential Research Program Project 2: Transparent Multimodal Crew Interface Designs Technical Note 3: Multimodal Cueing for Transparency in Mobility Operations*. 49.

Draper, M., Calhoun, G., Hansen, M., Douglass, S., Spriggs, S., Patzek, M., Rowe, A., Evans, D., Ruff, H., Behymer, K., Howard, M., Bearden, G., & Frost, E. (2017). Intelligent Multi-Unmanned Vehicle Planner with Adaptive Collaborative/Control Technologies (Impact). *19th International Symposium on Aviation Psychology*, 226–231.

Endsley, M. R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *37*(1), 32–64. https://doi.org/10.1518/001872095779049543

Hastie, H., Chiyah Garcia, F. J., Robb, D. A., Laskov, A., & Patron, P. (2018). MIRIAM: A Multimodal Interface for Explaining the Reasoning Behind Actions of Remote Autonomous Systems. *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, 557–558. https://doi.org/10.1145/3242969.3266297

Jiun-Yin Jian, Ann M. Bisantz, Colin G. Drury, & James Llinas. (1998). *Foundations for an Emprically Determined Scale of Trust in Automated Systems*. U.S. Air Force Research Laboratory.

Kaplan, A. D., Kessler, T. T., Brill, J. C., & Hancock, P. A. (2021). Trust in Artificial Intelligence: Meta-Analytic Findings. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 001872082110139. https://doi.org/10.1177/00187208211013988

Knight, W. (2021, May 10). The Pentagon Inches Toward Letting AI Control Weapons. *Wired*. https://www.wired.com/story/pentagon-inches-toward-letting-ai-control-weapons/

Lakhmani, S., Abich, J., Barber, D., & Chen, J. Y. C. (2016). A Proposed Approach for Determining the Influence of Multimodal Robot-of-Human Transparency Information on Human-Agent Teams. In D. D. Schmorrow & C. M. Fidopiastis (Eds.), *Foundations of Augmented Cognition: Neuroergonomics and Operational Neuroscience* (Vol. 9744, pp. 296–307). Springer International Publishing. https://doi.org/10.1007/978-3-319-39952-2_29

Lee, J. D., & See, K. A. (2004). Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, *46*(1), 50–80.

Lyons, J. B. (2013). Being transparent about transparency: A model for human-robot interaction. *2013 AAAI Spring Symposium Series*, 48–53. http://www.aaai.org/ocs/index.php/SSS/SSS13/paper/view/5712

Lyons, J. B., Koltai, K. S., Ho, N. T., Johnson, W. B., Smith, D. E., & Shively, R. J. (2016). Engineering Trust in Complex Automated Systems. *Ergonomics in Design: The Quarterly of Human Factors Applications*, *24*(1), 13–17. https://doi.org/10.1177/1064804615611272

Lyons, J. B., Sadler, G. G., Koltai, K., Battiste, H., Ho, N. T., Hoffmann, L. C., Smith, D., Johnson, W., & Shively, R. (2017). Shaping Trust Through Transparent Design: Theoretical and Experimental Guidelines. In P. Savage-Knepshield & J. Chen (Eds.), *Advances in Human Factors in Robots and Unmanned Systems* (pp. 127–136). Springer International Publishing.

Mercado, J. E., Rupp, M. A., Chen, J. Y. C., Barnes, M. J., Barber, D., & Procci, K. (2016). Intelligent Agent Transparency in Human–Agent Teaming for Multi-UxV Management. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *58*(3), 401–415. https://doi.org/10.1177/0018720815621206

Miller, C. A. (2021). Trust, transparency, explanation, and planning: Why we need a lifecycle perspective on human-automation interaction. In C. S. Nam & J. B. Lyons (Eds.), *Trust in Human-Robot Interaction* (pp. 233–257). Elsevier. https://www.sciencedirect.com/science/article/pii/B9780128194720000113

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *39*(2), 230–253.

Parasuraman, R., & Wickens, C. D. (2008). Humans: Still Vital After All These Years of Automation. *Human Factors*, *50*(3), 511–520. https://doi.org/10.1518/001872008X312198

Rafael, S. (2020). Multimodality, Naturalness and Transparency in Affective Computing for HCI. In A. Marcus & E. Rosenzweig (Eds.), *Design, User Experience, and Usability. Interaction Design* (pp. 521–531). Springer International Publishing.

Sheridan, T. B. (1988). Trustworthiness of Command and Control Systems. *IFAC-PapersOnLine*, *21*(5), 427–431. Supplemental Index.

Taylor, G., Purman, B., Schermerhorn, P., Garcia-Sampedro, G., Hubal, R., Crabtree, K., Rowe, A., & Spriggs, S. (2015). Multi-modal interaction for UAS control. *Unmanned Systems Technology XVII*, *9468*, 946802. https://doi.org/10.1117/12.2180020

Theodorou, A., Wortham, R. H., & Bryson, J. J. (2017). Designing and implementing transparency for real time inspection of autonomous robots. *Connection Science*, *29*(3), 230–241. https://doi.org/10.1080/09540091.2017.1310182

Vered, M., Howe, P., Miller, T., Sonenberg, L., & Velloso, E. (2020). Demand-Driven Transparency for Monitoring Intelligent Agents. *IEEE Transactions on Human-Machine Systems*, *50*(3), 264–275. https://doi.org/10.1109/THMS.2020.2988859