Building Trust in Human-Machine Teaming for Autonomous Space Sensing

Garrett Fitzgerald
United States Space Force

Rachel Morris
Applied Decision Science, LLC, Dayton, OH

Laura G. Militello
Applied Decision Science, LLC, Dayton, OH

Justin Fletcher
United States Air Force

ABSTRACT

Trust is a critical acceptance criterion for military applications of autonomy. Conflict, including conflict that extends into the space domain, involves complex and dynamic scenarios through which human and autonomous actors must jointly reason and act. Designing autonomous systems that are capable of earning operator trust while performing autonomous sensor orchestration in support of Space Domain Awareness (SDA) has proven particularly challenging, in part, due to the limited number of subject matter experts (SMEs) available for systematic survey and the absence of a framework by which to assess operator satisfaction with system outcomes. In this work, we employ Cognitive Task Analysis (CTA) to elicit and contextualize SME feedback, and suggest ways to incorporate this feedback during human-machine teaming interface development. We establish a tailored framework to explore SME trust in a ground-based telescope autonomous sensor orchestration system, MACHINA, using prior concepts in Human-Machine Teaming (HMT) and cognitive engineering. A HMT Knowledge Audit is adopted for the SDA telescope orchestration task and employed as a knowledge elicitation vehicle for United States Space Force SDA operators. Analysis of interview data from three user groups is mapped to trust antecedents, yielding actionable HMT guidance for system development. This work presents examples of how SME trust in AI can be improved throughout the development of autonomous space sensing applications by combining human factors engineering and HMT concepts.

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval LeMay Center/PA on 18 Dec 23; Case number AETC-2023-1317

1. INTRODUCTION

Since the dawn of the Cold War, the function of Human-Machine Teaming (HMT) under existential risk has preoccupied the imaginations of general officers and the general public alike. Catastrophe stemming from human-autonomy pairing gone wrong has been a pervasive idea in fiction, perhaps first espoused with *Dr. Strangelove's* doomsday machine. These ideas have been recently formalized in cautionary analyses which warn against heedless implementation of AI-enabled military decision support tools [1, 2, 3]. Although caution is warranted, inaction also imposes risk. To bridge this gap, designers of autonomous systems must actively cultivate trustworthiness.

Applications of recent advances in AI have the potential to augment virtually every mission and function of the Department of the Air Force. In response to the growth in both population [4] and dynamism [5] of Resident Space Objects (RSOs), the United States Space Force (USSF) has fielded a prototype autonomous system to responsively orchestrate Space Domain Awareness (SDA) sensors [6]. The work of expert operators in a dynamic environment can be aided by decision-support tools in planning observations, managing sensor availability across organizations, and evaluating observation results. However, the potential advantages of automation will remain unrealized if insufficient care is taken to develop a system that is both performant and trustworthy.

In this work, we identify leverage points for improving human-machine teaming in the context of SDA. Our research team interviewed future users of an autonomous space sensor orchestration solution, MACHINA, to better understand the cognitive challenges users manage today and implications for HMT. Findings are characterized in terms of McDermott et al's *HMT Knowledge Audit* [7] and antecedents of trust adapted from Mayor et al's seminal work on human-human trust [8].

1.1 Space Domain Awareness

The USSF SDA core competency entails continuous monitoring of RSOs in all orbital regimes within the gravitational influence of Earth [9]. This includes the detection and tracking of objects, object characterization, and data integration and exploitation from a global sensor network composed of radars, optical telescopes, and space-based assets. Each RSO must be continuously accounted for with a known orbit, updated with periodic observations to maintain an object catalog. In addition to scheduling updated orbital measurements for known objects, the system must also account for the identification of new objects. For deep-space RSOs (i.e., medium earth orbit (MEO) to geosynchronous equatorial orbit (GEO)), heterogeneous narrow field-of-view electro-optical (EO) sensors are primarily employed, necessitating complex orchestration functions for tasking and scheduling. With competing priorities for scheduling observations of target objects, the USSF network of globally dispersed EO sensors requires a dynamic orchestration system driven by automation.

Although various organizations are devoted to supporting different levels of scope within the overall strategy of SDA, orchestration tasks can be binned into general functions regardless of whether a unit is controlling a single telescope asset or managing a global network. Ground-based EO sensors perform the following functions to accomplish situational awareness (SA) for RSOs:

1. *Object Custody* includes periodic observation of an object. This is a system level capability using contributions from all available sensors.
2. *Object Proximity Monitoring* refers to the management of indications and warnings to support space asset protection. This is carried out by searching the area around an object of interest, reporting discovered objects, and maintaining custody of those objects.
3. *Responsive Object Monitoring* includes custody maintained through maneuvers and anomalous activity.
4. *Object Characterization* refers to measurable characteristics of an object such as size, rotation periodicity, and brightness pattern.
5. *Uncued Wide-Area Search* refers to scanning for space objects not yet accounted for.

Effective utilization of a complex global network of sensors and related orchestration is an ongoing challenge for the USSF. Successful SDA is contingent on the tasking efficiency and bandwidth of ground-based EO telescope systems. The large volume occupied by RSOs in various orbits (a search space of 73 trillion cubic miles from Low Earth Orbit

(LEO) to GEO), the increased proliferation of RSOs as illustrated in Figure 1, combined with the limited number of ground-based EO sensors creates significant operational challenges. Although the ability of EO systems to discover, track, and characterize objects in Earth's orbit allows for the effective employment and operation of satellites in all orbits, current systems will struggle to keep pace with the increasingly congested and complex space environment of the future [4, 10].

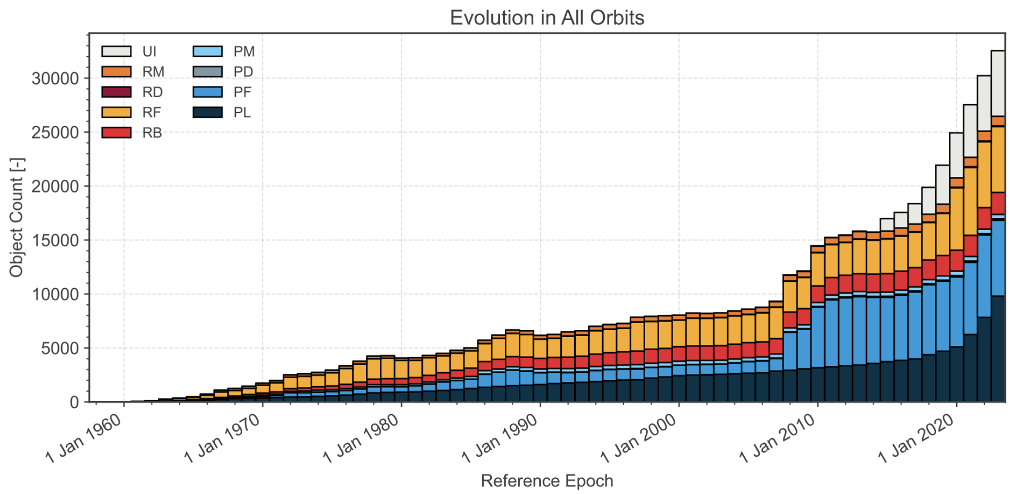

Fig. 1: The population of RSOs increased rapidly from 1960 to 2023, and is expected to continue to grow in the future. [11]

In order to effectively schedule and task a global system of EO telescopes under changing conditions, a flexible objective decomposition approach is needed [12]. There exist various frameworks for autonomous sensor tasking systems and learned reward algorithms for tasking [13, 14, 15], but producing a globally optimized schedule of sensor observations is NP-hard, severely complicating real-time dynamic optimal sensor scheduling [16, 17]. There is no single metric to quantify optimal command and control scheduling performance in the dynamic environment of SDA, although recent work has made strides to evaluate reward functions for schedule optimization [18].

Assuming a heuristic optimization approach is implemented in an autonomous orchestration system, this work in turn focuses on the human-machine teaming aspect of deploying a system which enables automation in scheduling and tasking of EO assets. This novel study into trust in autonomously managed SDA mission functions presents generalized cognitive requirements which will improve system HMT.

1.2 MACHINA

MACHINA (**M**ission driven, **A**utonomous, **C**ollaborative, **H**eterogeneous, **I**ntelligent, **N**etwork **A**rchitecture) is an autonomous SDA sensor orchestration solution, enabling operators to provide *Objectives* which are then decomposed into *intents* that are, in turn, translated to collection requests for sensors. MACHINA orchestrates a global network of EO sensors and is implemented as an agent-based software framework [6]. Delegating low-level management of a distributed network of sensors to MACHINA agents enables shared situational awareness for operators and AI to provide responsive SDA sensor orchestration while reducing operator workload. Under this framework human input supplies high-level objectives for multiple sensors (such as Object Custody for an object in GEO orbit), with sensor orchestration handled by MACHINA agent logic.

MACHINA operation begins with a set of *Objectives* and *Intents*. As the goal for MACHINA operation is multi-agent autonomy, the input of *Objectives* or *Intents* to scheduled *Collect Requests* can be prompted by users or machine-machine tasking. *Objectives* are higher order mission-level functions such as Catalog Maintenance or Object Search. Input variables for *Objectives* are general parameters such as start and end time, and priority level. *Intents* are instruction-level parameters for generating *Collect Requests*, including target and constraints. *Collect Requests* are target and sensor specific. *Objectives* may generate multiple *Intents*, which in turn generate *Collect Requests*, sent to the Scheduler. MACHINA's Scheduler aims to optimize the performance of EO collection, as measured by a configurable reward function, while remaining subject to the constraints implied by the objectives provided. MACHINA logic is illustrated in Figure 2.

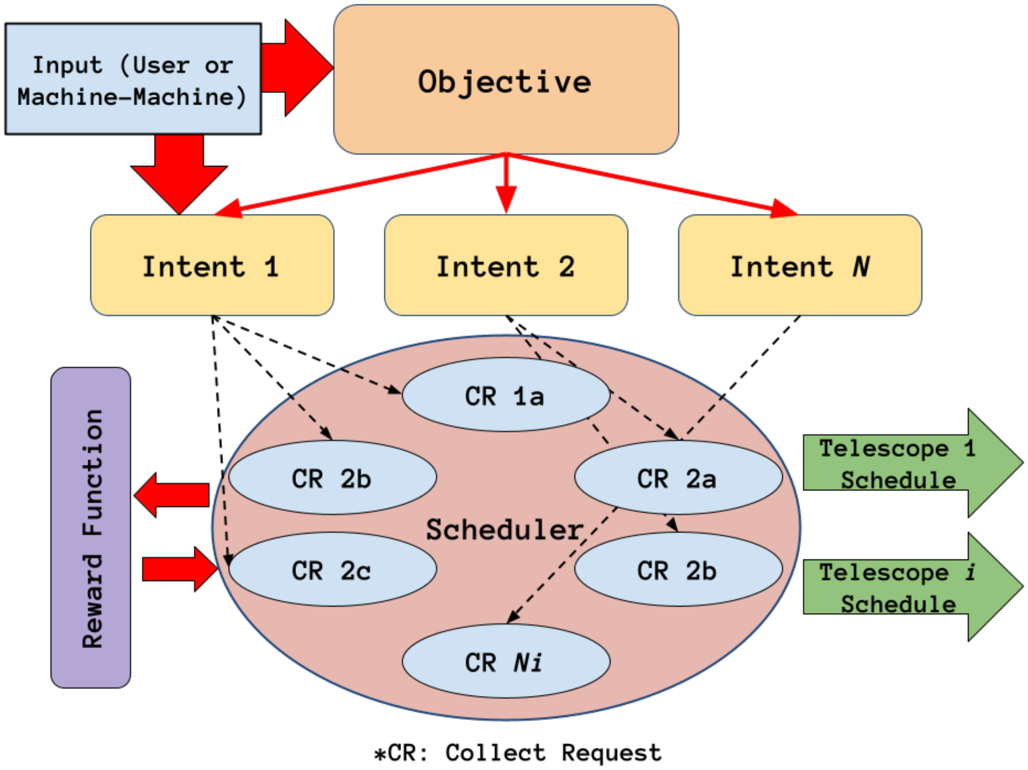

Fig. 2: An Illustration of the MACHINA agent logic workflow for scheduling. Humans input objectives, which are decomposed into intents. Intents are a collection of constraints that are provided to a scheduler. The scheduler produces a planned schedule comprising collection requests. These are, in turn, transmitted to sensors for collaborative fulfillment.

At present, MACHINA is in its beta development phase, with continuous additions being made to the system user interface (UI) and Agent logic. This timing presents a perfect opportunity to study trust in HMT for MACHINA employment, as important features to enable trustworthy operation can be discovered and implemented before system acceptance. Beta prototypes allow developers to begin to test their assumptions and determine which features should be built and made visible to users. User feedback at this phase can inform the UI and functionality for existing features, as well as envisioned capabilities (e.g., the addition of a Large Language Model for interfacing with MACHINA).

1.3 Exploring Trust in MACHINA Human-Machine Teaming

As technology evolves and the roles of humans and machines change, the language to describe these human-machine interactions also changes. Early references to man-machine systems have been replaced by more recent labels such as Human-AI teams, complementary intelligence, Human-computer interaction, and collaborative intelligence. For the purposes of this article, we consider all of these under the umbrella of HMT.

Many studies in HMT are twofold in approach to methodology: they define a framework or taxonomy for human-AI partnership with which they can measure system effectiveness, and then employ interview or survey techniques to measure user feedback relative to their definitions of HMT design themes [19]. Artifacts from HMT studies can uncover successful design traits and specific features which enable enhanced human-AI partnership. For example, an online survey found that when using an AI assistant for various complementary exercises such as object identification, user trust was largely predicated on clear communication of the AI confidence level to the human throughout tasks [20]. The literature on HMT is rich with similar studies which yield design improvements for human-AI systems. Trust is a critical facet of effective human-AI teaming. Humans must be able to accurately calibrate under what circumstances they can rely on technology and recognize when the technology is reaching the edges of its capabilities. Much of the literature on trust in HMT emphasizes design implications, such as transparency, communication of intents, and resilience as important contributors to AI trust and trustworthiness [21, 22, 23]. One study, conducted in 2023, illustrates the negative effect on human perceptions of trust and trustworthiness following an unexpected action or behavior from the AI system [24].

For this work, we used the HMT knowledge audit developed by McDermott et al based on an integration of the human-machine teaming literature [25, 7]. The HMT Knowledge Audit is a cognitive task analysis (CTA) technique tailored to elicit examples of effective HMT and highlight opportunities to improve trust and teaming. CTA is a core set of methods within cognitive engineering, used to model the decision making process used by SME operators as they perform their work [26]. It is used to systematically categorize, connect, and prioritize SME decisions in order to understand successful strategies used to accomplish mission objectives [27]. CTA provides a means of gaining meaningful insights into pressure points of what trust means by conducting in-depth interviews with experienced practitioners.

In parallel, we adapted the Ability, Benevolence, Integrity (ABI) framework of human-human trust [8] to characterize features of MACHINA (existing, planned, and envisioned) that improve or degrade perceptions of trust in the system. Specifically, we recast the definitions to better represent aspects of trust in HMT, and added the concept of task assignment. Task assignment was not derived from the literature on human-human trust; rather, it emerged as important component of trust in studies of HMT [24]. Task assignment refers specifically to tasks that culturally or ethically cannot be accomplished by technology without human approval. Table 1 summarizes our adaptations.

Antecedent of Trust	**Mayer et al (1995) Definition**	**Definition Applied to HMT**
Ability	The group of skills, competencies, and characteristics that enable a party to have influence within some specific domain.	Can the software do what I need it to do? (Calculation-based competence)
Benevolence	The extent to which a trustee is believed to want to do good to the trustor.	Do the software and human have shared SA and support collaborative joint cognition? (Mission-based competence).
Integrity	The trustor's perception that the trustee adheres to a set of principles.	Is technology performance consistent with human expectations?
Task Assignment		Are humans willing to hand this task over to AI?

Table 1: Definitions of Trust Antecedents

We applied a systems engineering approach to building trust in the HMT development process by using cognitive engineering methods to elicit and analyze "trustworthy" design elements related to HMT themes. We used CTA to systematically identify, connect, and prioritize cognitive challenges described by SMEs. These informed our understanding of successful strategies for accomplishing mission objectives, and how these strategies can be evolved in tandem with an AI partner. CTA methods are particularly useful for small sample in-depth investigations. In the context of the current study, CTA provides a means for eliciting operator expertise and insights to characterize critical enablers and inhibitors of trust in advanced automation for SDA. Using HMT design themes as a framework for crafting CTA interviews, and working with SME operators, we analyzed the HMT themes that are expressly important in trustworthy teaming with AI. In this paper, we offer examples of cognitive requirements related to HMT themes, which can be used to improve MACHINA as well as future AI systems in related USSF mission areas.

2. METHODS

We describe how the CTA process was applied to our study to extract operator insight on how SDA sensor orchestration functions are performed, including knowledge elicitation and data analysis.

2.1 Participants

We conducted interviews with three user groups including MACHINA user-developers, 18th Space Defense Squadron Det 1, and GEODSS telescope operators.

MACHINA User-Developers The User-Developers test new features on 15th Space Surveillance Squadron sensors in Hawaii, Colorado, and other locations. They ensure that newly implemented functional requirements for MACHINA work as intended. They perform routine collections on RSOs to calibrate and validate system end-to-end performance, measuring performance metrics such as astrometric residuals, uncorrelated target (UCT) rate, incorrect correlation rate, object detection model performance, and *Intent* failure reasons. They use all functions available on MACHINA, as the validation arm of development.

18th SDS Det 1 This user group performs actionable SDA functions on National Space Defense Center (NSDC) designated targets. Det 1 is threat focused and will maintain SDA relevant to known threats as well as identifying potential emerging threats to high-priority US and allied assets. A dynamic and time sensitive mission, operators would use MACHINA at the *Objective* and *Intent* level.

GEODSS The GEODSS system is a global network of EO telescopes used to monitor deep-space objects (10km - 45km) from earth (generally speaking, MEO to GEO orbits). GEODSS observations contribute to the Space Surveillance Network (SSN), fulfilling numerous SDA functions from catalog maintenance to wide-area search. Outside of day-to-day operations, specific taskings may come from different government agencies, which require shifting to collect data on high priority objects. Operators may input MACHINA *Objectives* or *Intents* and will be able to prompt *Objectives* which schedule the global network of sensors autonomously.

We interviewed 3 MACHINA user-developers, 4 future users from 18th SDS Det 1, and 1 GEODSS operator. Table 2 details the end user group of each interviewee and their years demonstrating the range of expertise associated with SDA.

End User Group	Participant	Experience
KBR User-Devs	1	9 Months
	2	4 Years
	3	1 Year
18 SDS Det 1	4	4 Years
	5	3 Years
	6	14 Years
GEODSS	7	40 Years

Table 2: Interviewee Demographics

2.2 Interviews

The interviews conducted as part of this study closely match the systems engineering approach outlined in [25, 7], with appropriate customization to the different mission areas of current and future MACHINA users. First, we reviewed HMT design themes. The HMT probe categories and related HMT themes offered by McDermott et all are detailed in Table 3. Their derivation is detailed in [7]. Each theme is the distillation of HMT research across multiple studies and systems. Using this framework, we created CTA knowledge audit probes tailored for the mission area of each of the three user groups.

This HMT Knowledge Audit follows the outline below, delving into individual SME decision making processes and perspectives on teaming with AI:

1. Baseline Mission and Top Challenges: Gain background understanding of mission area and role, elicit high-level cognitively and technically challenging job tasks.
2. AI Introduction: Establish rapport on MACHINA functionality and vision.
3. Critical Decision Method Probe: Unpack experiences which caused SME to improvise decisions outside of normal operations [28].
4. HMT Knowledge Audit: In depth CTA interview to probe SME experiences and preferences related to HMT themes [27].

Probe Categories	**Relevant HMT Themes**
Past and Future: Predictability, Exploring the Solution Space	Future intentions and activities are observable and understandable.
Big Picture: Observability	Transparency into what an AI partner is doing relative to the task.
Anomalies: Calibrated Trust, Directing Attention, Adaptability	Format information to support visualization of trends and anomalies
Noticing: Directing Attention, Information Presentation	Orient attention to critical problem features and cues.
Self-Monitoring: Common Ground, Calibrated Trust	Pertinent beliefs, assumptions, and intentions are shared.
Job Smarts (working efficiently)	Automation supports humans in working effectively.
Past and Future: Predictability, Exploring the Solution Space	Recognize and adapt fluidly to unexpected situations.

Table 3: HMT Probe Categories and Related Themes

A sample of the prepared interview guide questions for the HMT Knowledge Audit is itemized below. These specific probe questions are used to elicit feedback from 18 SDS Det 1 interviewees for the HMT theme category of *Past and Future: Predictability, Exploring the Solution Space*. Responses from each interviewee are be mapped to new cognitive requirements and trust antecedents.

- Q: Is part of your job to anticipate what might happen next? Can you give me an example? *Probe about existing tools and envisioned support*.
- Q: What historical information do you need to anticipate what will happen next? (e.g., do you need to understand typical satellite movements to predict future movements or spot anomalies?) Will you walk us through an example of how you use historical information? *Probe about existing tools and envisioned support*.
- Q: How predictable are your missions? Is one very much like any other, or do they vary in important ways?
- Q: What other kinds of variability do the users manage? What parts of this task are unpredictable?
- Q: Do you ever set tripwires or automated alerts? What sorts of things are you tracking in this way?

Sixty-minute interviews were conducted remotely via Microsoft Teams or face-to-face. Three researchers were present for each interview. One served as lead interviewer, others took notes and offered follow-up questions. Microsoft Teams interviews were recorded with the permission of interviewees. Interviews with representatives from the 18th SDS Det 1 were conducted in a facility that did not permit recording, so the team relied on handwritten field notes for those sessions. In addition to interviews with three individuals, a group interview was held with four representatives from the 18th SDS Det 1.

2.3 Analysis

The focus of our initial analyses was on identifying cognitive requirements. Cognitive requirements are specific cognitive activities or challenges that must be accomplished by the HMT; they generally represent important opportunities to improve teaming and facilitate trust. At least two members of the research team independently reviewed notes from each interview to identify cognitive requirements and link them to one or more HMT themes illustrated in Table 3. When cognitive requirements had been identified, all three team members independently categorized them according the trust antecedents. Follow-up meetings with all members of the research team were held to reach consensus on the trust antecedents. This process yielded 22 operator insights that were applicable to at least one HMT category and had important implications for trust.

3. RESULTS

We identified cognitive requirements linked to HMT themes, and explored the implications of cognitive requirements for trust. We describe each set of results in turn.

3.1 Cognitive Requirements and HMT Themes

We identified 22 cognitive requirements critical to effective HMT in the context of MACHINA. Presentation of all 22 is beyond the scope of this proceedings paper; instead we present a subset of illustrative examples. Table 4 includes an example of a cognitive requirement related to each HMT theme identified by McDermott et al [17]. In the far-right columns of Table 4, we indicate which antecedents of trust are supported by each requirement.

HMT Theme	**Cognitive Requirement**	**Example (Quote)**	**A**	**B**	**I**	**T**
Past and Future: Predictability, Exploring the Solution Space	Need ability to what-if potential courses of action: Dynamic Scheduler View.	"A preview button or predict button to display tentative schedule and play it forward would be very useful because 75 of those fails occur for a variety of reasons not obvious to the users."		✓	✓	
Big Picture: Observability	Need to know MACHINA's rationale behind the schedule	"When MACHINA develops a global schedule, users will need to know where that schedule comes from, what objectives are running the schedule."		✓	✓	
Anomalies: Calibrated Trust, Directing Attention, Adaptability	Need historical data to identify anomalies and inform revisit rate	"Operators are going to want to know the past maneuvers for an object. When was the last time I got an observation of this object? This helps minimize the revisit rate of a specific object."	✓	✓		
Noticing: Directing Attention, Information Presentation	Need tools to support noticing true maneuvers in the context of noisy data	"Want to have a maneuver detection capability... Automation needs to rapidly support custody and change detection...The current system would be fooled if two objects changed places during an observation gap."	✓			
Self-Monitoring: Common Ground, Calibrated Trust	Need to coordinate across sites to support common ground	"A future version of MACHINA might allow an operator in Maui to see what an operator in Colorado is scheduling. There might be a chat feature that would allow operators at different sites to chat with each other."		✓		
Job Smarts (working efficiently)	Need AI support minimizing revisit rates; achieve higher confidence with fewer observations	"Want AI to shorten timelines and reduce the number of required observations."	✓			
Improvising: Adaptability, Directability	Humans need the ability to assign priority	"I don't want MACHINA generating priority itself, I want to articulate to the automation what is important."		✓		✓

Table 4: Example Cognitive Requirements and Related HMT Themes with ABIT Assignment Categories

Past and Future: Predictability, Exploring the Solution Space Predictable automation may support an operator in exploring the solution space and identifying or predicting the future state of an object. One operator emphasized the value of a dynamic scheduler that would allow them to "what-if" or simulate potential courses of action. Other operators expressed a need for historical data on an object to understand its behavior as well as predict which sensors can achieve future collections. It is important that the MACHINA system supports the operator in accessing and

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval LeMay Center/PA on 18 Dec 23; Case number AETC-2023-1317

understanding past information and using it to "what-if" about the objects future state or behavior. These capabilities contribute to human trust and willingness to work with the automation by supporting collaborative joint cognition and expectations for how MACHINA will perform particular actions.

Big Picture: Observability Given the heavily congested environment in space and the competition for sensing resources, many operators advocated for a tool that presents information on a sensor's geographical locations, taskings associated with that sensor, and whether it is currently active on a collection. This information will provide the operator with a "big picture" view of the situation and the assets available for tasking. While many of the operators expressed interest in this capability, they also emphasized the importance of understanding the rationale, confidence, and risk or consequences associated with the proposed schedule. The interface should also allow the operator to see a detailed view of the schedule including the sensors to be tasked and external factors (e.g., scheduled maintenance, poor weather conditions) that may impact its collection capabilities. Clear representation of system priorities, actions, and rationale improve trust in automation by allowing the human to gauge if the technology's behavior is in line with their own goals.

Anomalies: Calibrated Trust, Directing Attention, Adaptability As operators work to maintain SDA, there is a need for tools to direct the human's attention to anomalies or unexpected situations. The operator must have the ability to quickly react to changing conditions at sensors, and this information should notify the operator in a manner which does not go overlooked. During one interview, the operator described a scenario in which there was a missed detection for a satellite. For this operator, the ideal response is for MACHINA to adaptively respond to the situation. These responses may include alerting the human to missed detections beyond operator-specified thresholds, and providing information about where the object might be. Providing historical data, supporting the human in investigating and adapting to anomalous events, and correlating detections to specific known satellites when applicable are useful design considerations for the MACHINA system and interface. Enabling the MACHINA system to maintain historical information on an RSO and notify the operator of inconsistencies demonstrates shared human-machine of goals and priorities, improving trust and trustworthiness in the system.

Noticing: Directing Attention, Information Presentation In addition to highlighting anomalies, it is also important that tools support users in noticing "typical" behaviors. These routine events and behaviors include maneuver detections and violated revisit rate thresholds that inform the operator how the mission is progressing. Many operators reported a need for a revisit rate metric that describes how a given sensor is performing based on current network limitations. These insights provide justification for two design implications. First, the MACHINA system should include a maneuver detection capability that details the objects' burn direction and proximity to other objects as well as suggesting changes to the scheduler to ensure collections on the object are still achieved. Second, the system should make visible the revisit rate for an RSO of interest as well as present the percentage of time a revisit rate is not met. These implications will allow the operator to respond proactively to object maneuvers as well as adjust taskings when revisit rates are not met and observation gaps are long. Consistent demonstration of competence and reliability from MACHINA are important factors that contribute to the humans trust in the system.

Self-Monitoring: Common Ground, Calibrated Trust Continuous improvement to advanced automation requires evolving trust calibration. The operator should understand the capabilities and limitations of the automation, which are likely to change over time as new sensors and data exploitation features are introduced. The automation should facilitate collaboration between human operators and AI so humans understand changing capabilities and limitations and can rely on MACHINA when appropriate and compensate for MACHINA limitations when needed. One interviewee articulated a need for support to coordinate scheduling with other geographically dispersed sensing sites. Currently, operators have little to no information on the collection requirements of other sensor locations. To address this, the MACHINA system should clearly display successful observations, objectives executed, and new opportunities for collections to support the operator in communicating with leadership and other operators. Additionally, MACHINA will need the ability to deconflict competing collection requests based on priority and future opportunities allowing the operator to focus on current collection activities. As space becomes more congested with RSOs, establishing common ground between human operators and MACHINA will be critical for effective HMT and calibrated trust in the automated assistant.

Job Smarts: Working Efficiently In some cases, SDA is characterized by high-stakes, time-pressured objectives often with geopolitical implications. As such, several operators reported the desire for automation that can shorten collection timelines and reduce the number of observations needed on an object to achieve high confidence in its position, capabilities, etc. They described the importance of minimizing revisit rates with fewer and shorter observations to allow more sensors to be tasked to a greater number of collect requests. As the development and design of MACHINA moves forward, it needs the capability to parse together information and imagery to create a comprehensive understanding of the target with a shorter stare time for the telescope systems. Features that supports the operator in accomplishing tasks more efficiently and accurately promote trust.

Improvising: Adaptability, Directability Humans are well-known to be highly adaptable and able to quickly improvise in response to an unexpected situation. Automation, however, tends to fail during these off-nominal situations requiring the human to manually control the system without the support of AI. Additionally, human operators may have information that the automation does not or cannot know. For example, one operator described the ability to input event-driven priority into MACHINA that articulates to the algorithm what is important. Another operator emphasized the importance of the human maintaining the final approval before a schedule is executed allowing them to make adjustments based on information external to the system. In this way, the human can trust that the system is considering their goals and priorities but will not execute an action without manual approval from the operator. Allowing the human to direct the system as needed by defining principles for the automation to adhere to is an important facet for improving human trust in MACHINA.

3.2 Cognitive Requirements and Trust Antecedents

Linking cognitive requirements to trust antecedents illustrates how cognitive requirements derived from HMT themes can promote trust. With regard to ability, interviewees offered examples of features that would give them confidence that MACHINA would support the tasks they need to accomplish to operate safely and effectively. For example, the need for historical data presented in a visualization that supports users in identifying anomalies is critical to SDA. A system that increases the ability of the human-machine team to do this better than either the human or the machine could do on their own would promote a high-level of trust. Similarly, a system that improves the ability to notice true maneuvers in the context of noisy data is likely to be enthusiastically embraced by users.

We used the term benevolence to refer to cognitive requirements related to high-level mission-based competencies. For example, interviewees described the need for a tool that supports coordination across sites to establish common ground in support of the larger Space Force mission; currently operators have little insight into the priorities and schedules of other sites. A technology that supports high level goals for the larger system is likely to engender trust.

Integrity refers to whether the technology's performance is consistent with user expectations. This is directly related to HMT themes of predictability and calibrated trust. One MACHINA-related example is the need for humans to know MACHINA's rationale behind the schedule. If users are to trust that the AI-generated schedule is actually accounting for priorities and constraints (i.e., space objects are only visible to specific sensors during specific time windows), they must have visibility into the rationale MACHINA uses to make scheduling decisions and deconflict priorities. A 'Dynamic Scheduler View' would also support integrity. Users would be able to play forward a schedule to see how their Objectives and Intents would likely play out, helping them develop accurate expectations about MACHINA's behavior.

Task assignment is narrower than ability, benevolence, and integrity, but critical to trust in HMT. In the context of MACHINA, interviewees indicated that they would not trust MACHINA to make priority decisions; humans must retain authority over priority decisions. This is because humans are likely to have information that MACHINA will not have such as an awareness of sociopolitical events and changes in commander's intent.

3.3 Recommendations for Supporting Trust in MACHINA

Our interviews suggest that MACHINA developers have worked closely with future users to develop an AI that supports HMT. They have considered the work of SDA operators and included features intended to increase their ability to achieve operator-level and enterprise-level goals. During this beta development phase, many of the user-interface elements are being specified. The cognitive requirements described in this paper can help guide developers in designing interface features that promote HMT themes such as predictability and observability. If the users do not know what to expect from MACHINA or what it is doing, it is difficult to effectively calibrate trust. As MACHINA provides users with more information than they have had access to before, it will be important to develop visualizations that direct users' attention to anomalies and trends. MACHINA has the potential to support collaboration across sites in ways that are not feasible today; well-designed representations and communication tools have the potential to improve efficiencies in important ways. Similarly, over time the AI will likely be able to analyze trends and identify opportunities to increase efficiencies such as requiring fewer observations to maintain confidence in an object's state. However, for these efficiencies to be realized, human operators must understand the rationale behind changes. Finally, it is critical that developers consider cultural and ethical issues that influence task assignment. In this report, we offer a small excerpt of the cognitive requirements identified.

3.4 Scheduling Visualization Concept

During the group interview with four representatives from the 18th SDS Det 1, participants drew their vision for a visualization to support scheduling on a whiteboard. This design concept was a springboard for discussion of many of the cognitive requirements described in Section 3.1. In this section we discuss the envisioned scheduling visualization. Figure 3 and Figure 4 provides our recreation of these mock-up, highlighting critical features and capabilities described during the group interview.

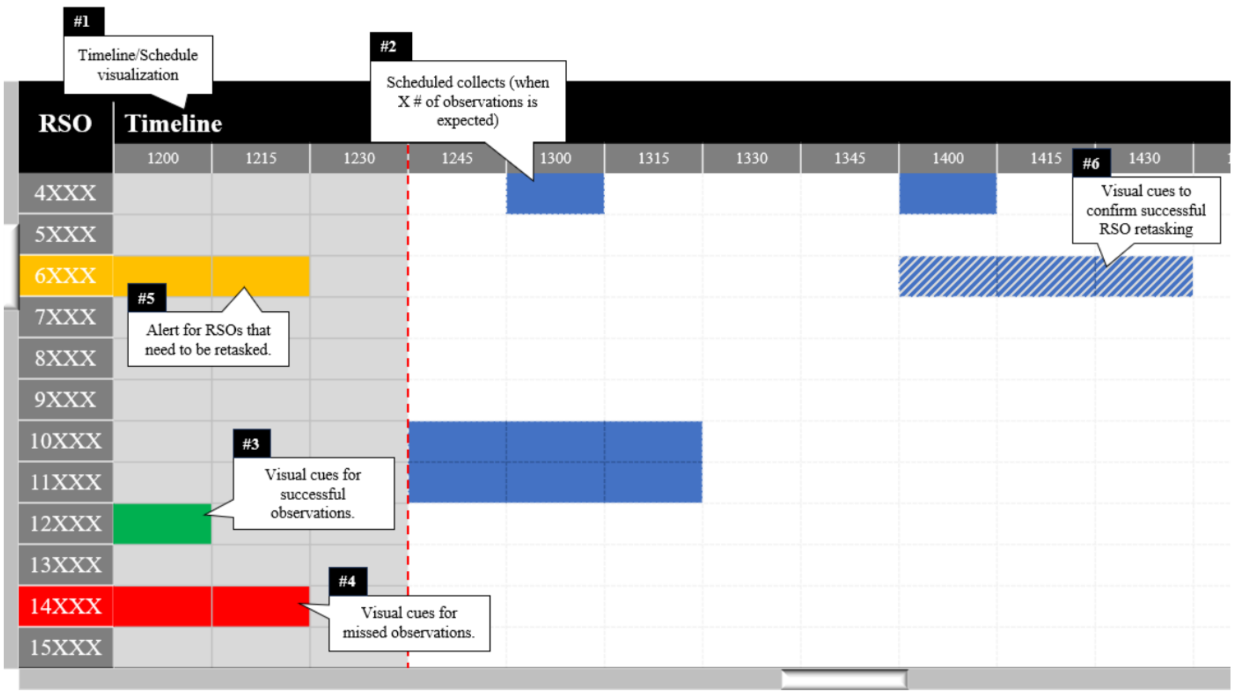

Fig. 3: User Interface Mockup A.

Timeline One key takeaway from this activity is the need for the operator to visualize and manipulate the schedule MACHINA generates. For example, one operator described the utility of looking into past and future collection requests to gauge the likelihood of mission success against variables that may be unknown to MACHINA. This feature is represented at the top of Figure 3 and Figure 4, in the form of a timeline (Figure 3, callout #1). This scrollable timeline would allow the users to navigate backward in time to view previous scheduled collections, and forward in time to see planned collections (Figure 3, callout #2). Historical observations that were met would be indicated in green (Figure 3, callout #3), and those that were missed would be indicated in red (Figure 3, callout #4).

This timeline view supports the operator in understanding and maintaining awareness of the RSO of interest as well as the window of time that observations of the RSO are expected. Additionally, the timeline view promotes trust by allowing the operator to look at previous as well as expected future data to evaluate and verify MACHINA's schedule. This view also establishes common ground as the operator can gain insight into what information MACHINA is using as well as input additional information unknown to the automation.

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval LeMay Center/PA on 18 Dec 23; Case number AETC-2023-1317

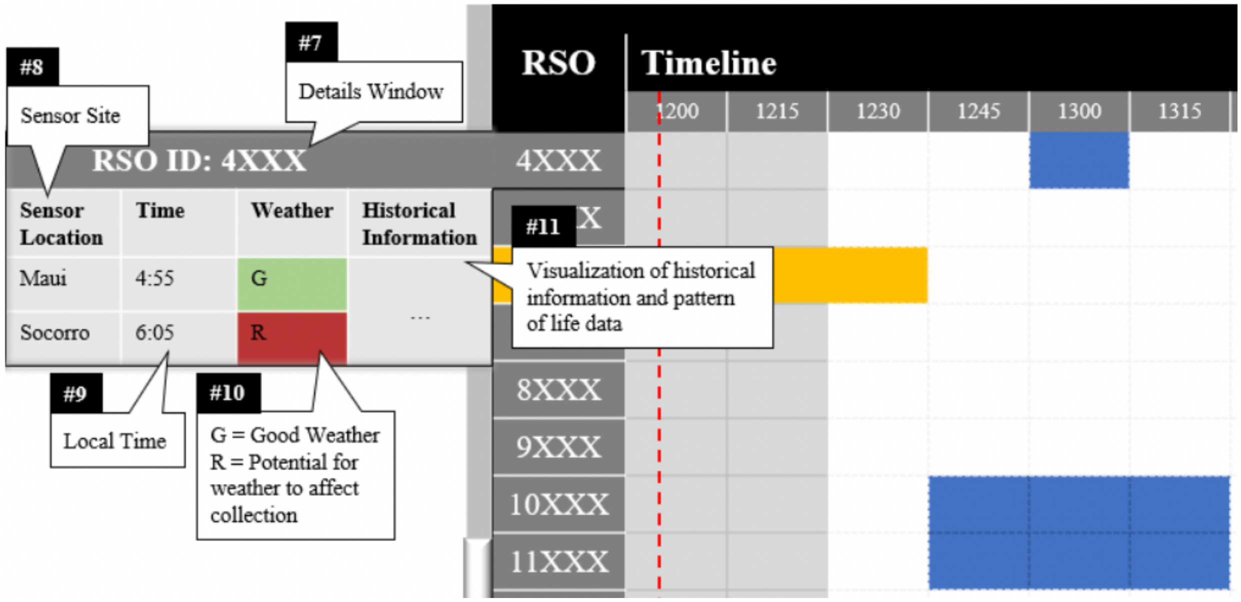

Fig. 4: User Interface Mockup B.

Alerts A second key takeaway from this exercise is the emphasis on alerts and the ability to create and modify them based on the mission and the target RSO. Operators described a need to set alerts for particularly critical RSO observations. Missed observations are routine and can often be deferred until the next scheduled observation, but in some cases more immediate retasking is needed. In this case the operator would have identified this particular RSO observation as critical. Therefore, the missed observation is highlighted in the schedule, calling the operator's attention to the need for retasking of a sensor (Figure 3, callout #5). The schedule would also indicate when a sensor has been tasked to observe the target RSO (Figure 3, callout #6). In addition, there may be value in setting an alert for RSOs that have missed several sequential scheduled collections. Enabling the ability to create and manage alerts allows the operator to focus on current tasks. This capability demonstrates the automation's ability to provide support on tasks that are important to the operator.

Details Window Our third takeaway is the need to access detailed information about RSOs. One design feature might be to make the RSO identifiers clickable; clicking on a specific RSO would invoke a details window (Figure 4, callout #7). The details window might include information such as the sensor site tasked for collection (Figure 4, callout #8), the local time at the sensor site (Figure 4, callout #9), and an indicator denoting weather conditions that the operator can use to gauge the potential for failed observations (Figure 4, callout #10). This window might also include historical information and pattern of life data on the target (Figure 4, callout #11).

This capability supports the operator in anticipating future problems and provides additional insight as to whether MACHINA's tentative schedule is feasible. This tab may also promote trust by making visible the capabilities of sensor sites as well as information relevant to MACHINA's scheduling decisions.

4. CONCLUSION

This study utilized an applied methodology for understanding the operators' role in SDA and generating cognitive requirements that enable trust and effective HMT with MACHINA. We performed a HMT Knowledge Audit on three end user groups of MACHINA, and analyzed the interview data to produce recommendations of system features which support trusted employment of the system. Across the DoD, performing this type of study throughout development of AI systems will ensure that operator feedback is addressed when adding autonomous functionality to systems. In this paper, we present a subset of the cognitive requirements identified in the CTA and an exploration of related antecedents of trust. We also outline a scheduling visualization design concept which demonstrates a system feature expressly supporting cognitive requirements and trust antecedents uncovered in the HMT Knowledge Audit. This foundational CTA can be used to aid designers in prioritizing user-interface design components and feature development to realize the potential of MACHINA for increasing the effectiveness and efficiency of adding autonomous functionality to SDA.

Recommended future work includes completing additional data collection and analysis with the other future user groups of MACHINA such as RAVEN and xGEO telescope operators and to further inform beta development. As MACHINA is implemented, it will be important to conduct follow-up interviews with each user group to identify failure points and areas for improvement.

Acknowledgements
The authors would like to thank the Space Systems Command MACHINA team for facilitating access to developers and user groups. GF expresses his gratitude to Laura Militello and Rachel Morris at Applied Decision Science, who are part of the world's best team in CTA. The authors also thank partners at MIT Lincoln Laboratory and the DAF AIA SDA project, including Jeffrey Price.

REFERENCES

[1] James Johnson. "Delegating strategic decision-making to machines: Dr. Strangelove Redux?" In: *Journal of Strategic Studies* 45.3 (2022). Publisher: Taylor & Francis, pp. 439–477.

[2] Ross Andersen. "Never Give Artificial Intelligence the Nuclear Codes". In: *The Atlantic* (May 2023).

[3] Cameron Hunter and Bleddyn E Bowen. "We'll never have a model of an AI major-general: Artificial Intelligence, command decisions, and kitsch visions of war". In: *Journal of Strategic Studies* 47.1 (2024). Publisher: Taylor & Francis, pp. 116–146.

[4] *Space Situational Awareness: DOD Should Evaluate How it can Use Commercial Data*. Tech. rep. GAO-23-105565. GAO, Apr. 2023.

[5] Clayton Swope. *Space Threat Assessment 2024*. Center for Strategic and International Studies (CSIS), 2024.

[6] Justin R Fletcher and Jonathan E Kadan. "Autonomous interactive agents for global telescope network orchestration". In: *Software and Cyberinfrastructure for Astronomy VIII*. Vol. 13101. SPIE, 2024, pp. 990–999.

[7] Patricia McDermott et al. "Human-machine teaming systems engineering guide". In: *MITRE CORP BEDFORD MA BEDFORD United States* (2018).

[8] Roger C Mayer, James H Davis, and F David Schoorman. "An integrative model of organizational trust". In: *Academy of management review* 20.3 (1995). Publisher: Academy of Management Briarcliff Manor, NY 10510, pp. 709–734.

[9] B. Chance Saltzman. *White Paper on Competitive Endurance: A Proposed Theory of Success for the U.S. Space Force*. Tech. rep. 2023.

[10] Mark R Ackermann et al. "A systematic examination of ground-based and space-based approaches to optical detection and tracking of satellites". In: 2015.

[11] ESA Space Debris Office. *European Space Agency's Annual Space Environment Report*. May 2023.

[12] Justin Fletcher. "The Dynamic Optical Telescope System: Collaborative Autonomous Sensing for Space Domain Awareness". In: *JDR&E* 4 (2021), pp. 10–18.

[13] Chenbao Xue et al. "Review of sensor tasking methods in Space Situational Awareness". In: *Progress in Aerospace Sciences* 147 (2024). Publisher: Elsevier, p. 101017.

[14] Cameron Harris et al. "Expanding the space surveillance network with space-based sensors using metaheuristic optimization techniques". In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Maui, HI, USA*. 2021.

[15] Robert Carden, Dustin Burchett, and Harvey Reed. "SNARE (sensor network autonomous resilient extensible): decentralized sensor tasking improves SDA tactical relevance". In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Maui, HI, USA*. 2021.

[16] Richard Linares and Roberto Furfaro. "An autonomous sensor tasking approach for large scale space object cataloging". In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*. 2017, pp. 1–17.

[17] Grant Nations and Justin Fletcher. “An Integer Programming Approach to Observation Scheduling for Space Domain Awareness”. In: *2024 IEEE Aerospace Conference*. IEEE, 2024, pp. 1–12.

[18] Jonathan Evan Kadan. “Reactive, Autonomous, Markovian Sensor Tasking in Communication Starved Environments”. In: (2024). Publisher: Virginia Tech.

[19] Erik Blasch, James Sung, and Tao Nguyen. “Multisource AI scorecard table for system evaluation”. In: *arXiv preprint arXiv:2102.03985* (2021).

[20] Qiaoning Zhang, Matthew L Lee, and Scott Carter. “You complete me: Human-ai teams and complementary expertise”. In: *Proceedings of the 2022 CHI conference on human factors in computing systems*. 2022, pp. 1–28.

[21] Klaus Christoffersen and David D Woods. “How to make automated systems team players”. In: *Advances in human performance and cognitive engineering research*. Emerald Group Publishing Limited, 2002, pp. 1–12.

[22] Sidney WA Dekker and David D Woods. “MABA-MABA or abracadabra? Progress on human–automation co-ordination”. In: *Cognition, Technology & Work* 4 (2002). Publisher: Springer, pp. 240–244.

[23] Glen Klien et al. “Ten challenges for making automation a ”team player” in joint human-agent activity”. In: *IEEE Intelligent Systems* 19.6 (2004). Publisher: IEEE, pp. 91–95.

[24] Joseph B Lyons, Izz aldin Hamdan, and Thy Q Vo. “Explanations and trust: What happens to trust when a robot partner does something unexpected?” In: *Computers in Human Behavior* 138 (2023). Publisher: Elsevier, p. 107473.

[25] Patricia L McDermotta et al. “Quenching the thirst for human-machine teaming guidance: Helping military systems acquisition leverage cognitive engineering research”. In: *13th International Conference on Naturalistic Decision Making*. 2017, pp. 236–240.

[26] Robert R Hoffman and Laura G Militello. *Perspectives on cognitive task analysis: Historical origins and modern communities of practice*. Psychology Press, 2008.

[27] Laura G Militello and Robert JB Hutton. “Applied cognitive task analysis (ACTA): a practitioner’s toolkit for understanding cognitive task demands”. In: *Ergonomics* 41.11 (1998). Publisher: Taylor & Francis, pp. 1618–1641.

[28] Robert R Hoffman, Beth Crandall, and Nigel Shadbolt. “Use of the critical decision method to elicit expert knowledge: A case study in the methodology of cognitive task analysis”. In: *Human factors* 40.2 (1998). Publisher: SAGE Publications Sage CA: Los Angeles, CA, pp. 254–276.