Trust in Autonomous Systems for Threat Analysis: A Simulation Methodology

Gerald Matthews[1(✉)], April Rose Panganiban[2], Rachel Bailey[2], and Jinchao Lin[1]

[1] Institute for Simulation and Training, University of Central Florida,
Orlando, FL, USA
{gmatthews, jlin}@ist.ucf.edu

[2] Air Force Research Laboratory, Wright-Patterson AFB, OH, USA
{april_rose.fallon, rachel.bailey.8.ctr}@us.af.mil

Abstract. Human operators will increasingly team with autonomous systems in military and security settings, for example, evaluation and analysis of threats. Determining whether humans are threatening is a particular challenge to which future autonomous systems may contribute. Optimal trust calibration is critical for mission success, but most trust research has addressed conventional automated systems of limited intelligence. This article identifies multiple factors that may influence trust in autonomous systems. Trust may be undermined by various sources of demand and uncertainty. These include the cognitive demands resulting from the complexity and unpredictability of the system, "social" demands resulting from the system's capacity to function as a team-member, and self-regulative demands associated with perceived threats to personal competence. It is proposed that existing gaps in trust research may be addressed using simulation methodologies. A simulated environment developed by the research team is described. It represents a "town-clearing" task in which the human operator teams with a robot that can be equipped with various sensors, and software for intelligent analysis of sensor data. The functionality of the simulator is illustrated, together with future research directions.

Keywords: Autonomous systems · Trust · Threat detection · Simulation
Cognitive processes

1 Introduction

1.1 Autonomy in the Military Context

The US military will increasingly rely on autonomous systems to perform actions currently delegated to human Warfighters, including detection of explosives, reconnaissance and surveillance, and support of combat operations. Such systems include robots and unmanned vehicles capable of independent situation analysis, decision-making and action, under some level of human monitoring and control. The US Air Force envisages autonomous systems making contributions to a range of operations [1]; we focus here especially on intelligence, surveillance, and reconnaissance (ISR). One realization of the Air Force vision is the "Loyal Wingman" concept, a

J. Y. C. Chen and G. Fragomeni (Eds.): VAMR 2018, LNCS 10910, pp. 341–353, 2018.
https://doi.org/10.1007/978-3-319-91584-5_27

scenario involving collaboration between a manned fighter platform and one or more Unmanned Autonomous Systems (UASs) with capabilities for locating and possibly attacking targets without full-time human direction. Autonomy may minimize cognitive load on the pilot, and protects the mission against jamming of communications between the pilot and the UAS.

Unmanned systems are especially suitable for "dull, dirty and dangerous" missions [2], including monitoring for threat. Consider, for example, a soldier at a vehicle checkpoint tasked with identifying possible insurgents. Use of a robot to detect hazardous materials such as explosives traces or radiation is within current capabilities. Advancements in machine intelligence will enhance robot functionality. For example, it might utilize infrared cameras to determine if the vehicle's body panels had been altered to hide contraband, a determination that requires complex inferences from sensor data. Robots will also acquire increasing abilities to analyze human beings for threat. Analysis of facial emotion and body posture can indicate fear and aggressive intentions, whereas off-the-body sensors will detect physiological responses such as autonomic arousal. For example, eye tracking methodologies show promise for detecting insider threat behavior at a computer workstation [3]. Effective use of such strategies requires more than advanced sensor technology. Psychophysiological responses do not map onto human behavior and emotion in a simple one-to-one manner [4]; they must be interpreted insightfully with some understanding of context, requiring "intelligence" on the part of the robot. For example, the Transportation Security Administration (TSA) has a pilot Behavior Detection and Analysis (BDA) program, which seeks to identify suspicious passenger behaviors at airports. However, behaviors alone may not be sufficiently diagnostic to be practically useful; analysis of the context for the behavior may be necessary to distinguish a fearful terrorist from a person afraid of flying.

An intelligent threat detection system would be of great value to the military and security services, through augmenting human capabilities, relieving human personnel of the tedious work of evaluating mostly harmless civilians, and physically removing humans from the potential dangers of close contact with those who are far from harmless. However, human oversight will remain critical. That is, threat detection and neutralization will require collaborative decision making between a human and an intelligent system. The robot or other autonomous agent must maintain and update threat evaluations based on its own sensors, communication with other team-members, and inference mechanisms.

1.2 The Importance of Trust

Teaming with autonomous systems places a burden of trust on the human operator [5]. Some degree of trust is essential to capitalize on the functionality of the autonomous system, but the operator must also remain alert to possible system errors, requiring careful calibration of trust. In the security context, human oversight is necessary to detect threats beyond the machine's detection capabilities, given the diversity of threats that may occur. It is also important to reduce "false positive" threat determinations which may waste resources and antagonize civilian populations.

There is a large human factors literature on trust [6, 7], which increasingly refers to human-robot interaction [8, 9]. Issues of automation misuse, disuse, and abuse [PR] broadly apply to autonomous as to other mechanical systems. A meta-analysis of trust in human-robot interaction [8] found that system performance characteristics including reliability, false alarm rate, and failure rate were more strongly related to trust than other robot attributes, or human and environmental characteristics.

Existing research provides only limited guidance for optimizing trust in autonomous systems operations [5]. A key issue is that enhancements in machine intelligence will change operator perceptions of functionality in complex ways, impacting trust in the process. On the one hand, greater intelligence will improve the machine's capabilities, enhance its ability to accommodate contextual factors, and improve its communications with the human operator. Generally, these capabilities should enhance trust. On the other hand, the downside of machine intelligence is that the bases for analysis and decision become increasingly complex and hard to communicate, and diagnosis of machine error is correspondingly difficult. Operators may also have faulty assumptions about machine intelligence that interfere with trust calibration. Thus, findings from trust research based on conventional automation may not generalize to autonomous systems.

1.3 Current Aims and Scope

Limitations of current research suggest a need for new methodological approaches to understanding the factors that impact operator trust in autonomous systems [5]. This paper describes some of the factors that require investigation, and describes a novel simulation methodology for determining how characteristics of the autonomous system, the operating environment, and operators themselves may influence trust. The methodology focuses on threat analysis as a specific context in which a human operator teams with an autonomous system. Specifically, the methodology aims to simulate interactions with an autonomous robot possessing sensors for various types of threat stimuli, as well as the capacity to analyze sensor data intelligently.

2 Trust in Autonomous Systems: Research Challenges

2.1 Facets of Trust

Trust is a complex construct with multiple facets and determinants [6]. Furthermore, it is a concept that comes from the social psychology of interpersonal relations, and its generalization to trust in machines is uncertain. In human factors research on trust in traditional automation, perceptions of competence predominate [9].

Qualities such as technical competence, reliability and understandability are readily applicable to machines [10]. Human-human trust also depends on additional factors - benevolence and integrity complement ability (competence) in one well-known model [11]. Such factors imply a self-motivated agent, perceptions that are unlikely to apply to simple automated systems, such as a vehicle cruise control. However, autonomous systems may be perceived as possessing a limited kind of personhood [12], implying that human-centered models of trust may become increasingly applicable.

Emotional as well as cognitive processes are also critical for human-human trust, which is influenced by the simple preference of liking or disliking the other person, a preference generated automatically with little cognitive effort [13]. Emotional aspects of trust are also shaped by longer-duration deliberative processes, e.g., a person might initially dislike a new coworker but come to appreciate their contributions over time. Simple automation elicits emotions associated with competence, such as frustration at repeated failures, or contentment with satisfactory performance. However, perceptions of machines as person-like agents open up the range of possible emotional responses, including socially-infused emotions. The person might experience disappointment or pride in the machine's performance, or guilt over failing to support it effectively.

2.2 Demand Factors and Trust

The advent of machine intelligence increases the range of demands potentially experienced by the operator, with implications for trust. Generally, as interacting with the machine becomes more demanding, trust is likely to deteriorate, as costs of machine management become perceived as higher than benefits of the machine's contributions to the mission [14]. Demands and trust may also be linked reciprocally; failure to calibrate trust optimally is likely to increase demands. Under-trust means that the human must take on more work, unnecessarily; over-trust will eventually lead to a mission failure due to machine error which the human must responsibility for mitigating.

Table 1. Challenges of autonomy and their performance impacts.

Challenge	Examples	Demands of autonomy	Performance vulnerabilities
Cognitive	• Sensor malfunction • Machine goal management error • Cyber attack	• Additional cognitive demand	• Attentional overload • Knowledge-based errors • Mode errors
Social	• Mis-perceptions of autonomy • Failure to support machine (human is poor team-mate) • Perceived lack of support (machine perceived as poor team-mate)	• Maintaining shared situation awareness • Appropriate back-up behavior • Function allocation	• Impaired shared situation awareness • Lack of team cohesion • Negative attitudes to machine
Self-regulation	• Stress from uncertainty and overload • Loss of perceived self-efficacy	• Overload • Managing feedback from machine (explicit or implicit)	• Diversion of attention • Disruption of executive control

Table 1 lists multiple sources of demand characteristic of interactions with autonomy [14]. The greater complexity of autonomous systems relative to traditional automation may impose higher demands on the operator. They may also increase operator uncertainty; for example, fault diagnosis becomes more difficult if it is unclear whether the fault is in a sensor or in the machine software. Demands may also be exacerbated by uncertainty over the machine' intentions [8]; in the military context, the operator might wonder whether unusual behavior was the result of the machine being hacked by a cyber-adversary. High cognitive demands and uncertainty may influence trust via the operator's assessment of machine competence; a machine that behaves unpredictably may not be deemed trustworthy.

Increasing autonomy also raises the scope for "social" interaction, as the machine graduates from tool to team-mate [15]. Team operations require not only coordination of actions to accomplish mission goals, but also teamwork behavior such as mutual performance monitoring, providing back-up, and leadership [16]. Lack of trust impairs teamwork, potentially leading to issues such as breakdown of a shared situational awareness and cohesion in performance [17]. Conversely, perceptions of poor teamwork by the machine – for example, if it fails to back-up the human as anticipated - will damage trust. Finally, interacting with an autonomous system may increase self-regulative demands, as the human is forced to evaluate their own competence as an operator. Increased cognitive and social complexity may make it difficult for the human to gauge if he or she is actually performing competently, potentially causing stress which may disrupt information-processing [14]. An officer commanding troops understands the importance of effective leadership – but what constitutes leadership of a team of autonomous systems? In some cases, the system may have the capacity to adapt its behavior according to its evaluation of the human's capabilities, i.e., adaptive automation [18]. The downside of this facet of machine intelligence is that the human may feel denigrated if the machine's actions signal that it perceives the human as incompetent.

3 Drivers of Trust in Autonomous Systems for Threat Analysis

Trust in autonomous systems may be influenced by novel demand factors, in addition to established drivers of trust associated with machine performance and reliability [9]. Salient demand factors will be somewhat context-dependent, varying with the functionality of the specific autonomous system, and the operational challenges faced by the human-system team. We discuss the factors that may be important in the threat detection and analysis context, which define research priorities.

3.1 System Characteristics

Simple devices for threat detection include sensors for radiation or toxic chemicals. Future systems will add to these capabilities in various ways including novel sensors such as lidar-based detection of threats in 3-D space. They will also include sensors to detect psychologically relevant human responses such as facial emotions and autonomic arousal, coupled with software that can distinguish, for example, harmless expressions of frustration from purposeful aggression. Human operators may find it hard to trust advanced and/or unfamiliar detection and analysis capabilities. It may be especially hard to trust a machine that makes psychological judgments.

Future autonomous systems will also differ from conventional automation in regard to communication abilities. Current systems provide essentially a passive read-out of information, and the human must gauge its credibility and whether it is actionable. Autonomous systems will be able to also deliver confidence ratings, background information and transparency messages [19] that illuminate the reasons for its analysis, potentially increasing cognitive demands. Future systems may also have action capabilities, potentially including autonomous search for possible threats, calling for human or machine back-up to deal with a threat, or in some cases direct actions such as bomb disarming. Such capabilities will increase "social" needs for effective teaming and appropriate trust calibration.

3.2 Environmental Characteristics

Threat detection in military and security contexts takes place in a variety of operating environments, differing in the challenges that they pose. In some cases, the environment may be generally safe, and threats rare, such as scanning people attending a sports event for traces of explosives. In other settings, the ability of the machine to improve over human threat detection may be critical. In the military context, threats may be hard to identify, due to the increasingly asymmetric nature of combat; for example, insurgents seek to blend in with civilians and change tactics frequently. Future threat identification will increasingly require information fusion [20], i.e., analysis of multiple cues provided by different information sources, often under time pressure. For example, the machine might analyze immediate emotional cues along with information about a suspect's social media postings, credit card purchases, and phone records. Multiple cue integration, perhaps including "big data" methods, may threaten transparency, even if such functionality is included; it may not be feasible to explain the machine's analysis to the human, placing a particular burden on trust.

3.3 Operator Characteristics

Human operators differ considerably in terms of the personal characteristics they bring to autonomous system interactions, such as level of understanding of information technology. Inexperienced operators may be vulnerable to misleading depictions of artificial intelligence in popular media. Various psychological factors associated with propensity to trust humans have been identified, but the extent to which they generalize to trusting autonomous systems is unclear, especially as the social behaviors that influence trust may be interpreted differently when executed by an artificial system.

Operator characteristics can be represented as mental models, in this context, the person's internal representation of the machine's capabilities and limitations [15]. In fact, two types of mental model are relevant. First, people will have pre-existing mental models of what artificial system capabilities. Barriers to trust include beliefs that machines cannot interpret human behaviors, or that machines cannot be relied upon as team-mates. Second, people will have a more narrowly focused model of machine-functioning in the context of the current mission's goals [17]. Training should accomplish realism in the mental model, but representations may nevertheless be biased by pre-existing conceptions of machine intelligence and by other characteristics such as personality. For example, in a UAS study, highly conscientious individuals appeared to prefer their own agency to reliance on automation under high demands [21].

3.4 Research Implications

We have identified multiple factors that may impact the operator's trust in an autonomous threat detection system, such as a robot or UAS. From previous research [8, 9] we can anticipate that factors such as perceived competence of the machine will impact trust, but system autonomy raises novel issues. Will the operator trust the machine's ability to infer threat from complex data sources, including psychophysiological data? Will the operator trust the machine to interpret threats emanating from challenging operational environments, such as insurgents actively aiming to disguise their intentions? Will the operator's trust be swayed by personal characteristics, such as pre-existing beliefs about intelligent machines, attitudes to technology, and knowledge of information technology?

The questions just framed point primarily to the role of cognitive demands in shaping trust, i.e., whether the operator can effectively cope with the increased complexity of working with an intelligent machine. However, in extended interaction scenarios, social and self-regulative demands may also factor into trust. Operators must determine the extent to which they will treat the machine as a team-mate capable of autonomous action, and the extent to which they trust their own judgment in managing the machine. Research on complex, realistic threat-detection scenarios may be necessary to answer such questions. Next, we describe a simulation platform that will be used in our research.

4 Simulation Use and Current Methodology

4.1 Unreal Engine for Research

Video game platforms are a favored tool for training purposes due to their versatility, elevating them to the status of "Serious games." [22]. Anecdotally, many military members report playing video games, supporting efforts to enhance training delivered through serious games. Younger Soldiers have greater gaming experience than older ones [23], so that gaming exposure is likely to increase in the future. The stimulating design of these systems may be naturally motivating for individuals [23] who regularly play video games. Serious games allow for the delivery of complex, situation-based information for the purpose of training due to the ability to program many environments and scenarios with multiple users. Using these features of serious games, researchers can examine decision-making, performance and subjective response in more realistic environments than traditional laboratory settings.

Gaming environments can be easily programmed for simulation of different complicated tasks from simple flight to dynamic exploration of a simulated environment. Unity and Unreal Engine 4 (UE4), by Epic Games, are popular game engines. The current study made use of UE4 due to its availability and ease of use. UE4 is a free downloadable gaming environment that is easily customizable. Objects for scenes can be purchased online and altered within the studio. Agents for simulation can be made in 3D modeling software such as make human (www.makehuman.org) or blender (www.blender.org) with animations created in software such as Mixamo (www.mixamo.com). Additionally, agents can be scanned into the designed environment and animated with commercial off-the-shelf software [24].

Scripting of levels can be done within UE4's editor using a node-based system called Blueprint shown in Fig. 1. Additionally, programming can be supplemented or done entirely using C++ and there is an abundance of tutorials for use of UE4 on YouTube and at the UE4 website (https://docs.unrealengine.com). The UE4 editor allows for programing of elements such as game rules, conditions, camera perspective, player control, weapon system controls, trigger events, and randomized or procedurally-generated props within the game [22]. Sound can also be added for atmosphere.

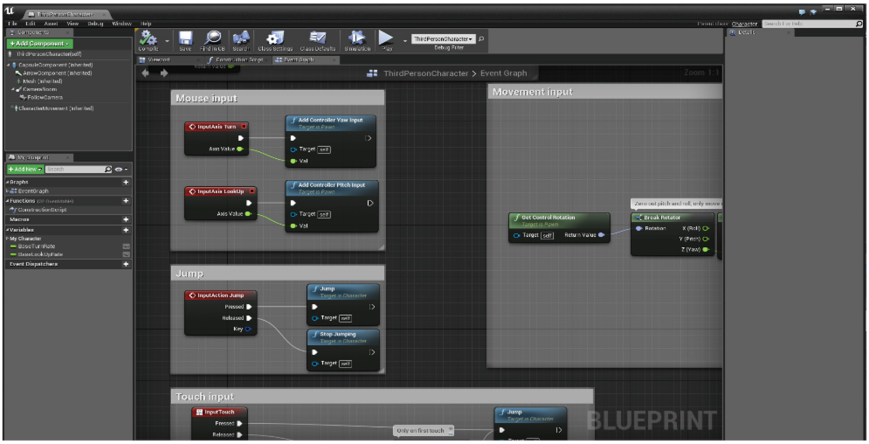

Fig. 1. Picture of UE4's node-scripting system.

4.2 Design of the Task

The current task was designed as a "town-clearing" task where a participant plays the role of a Soldier patrolling a small city with a robot partner to determine if threat activity is present. Participants are told that they are clearing a path through the town for a SWAT team to travel and must ensure that the areas along this path are safe. The robotic partner is armed with multiple sensors which it uses to make its own determination of the area's potential for threat. The type of cue used by the robot is manipulated. In one condition, the robot makes its judgement of the scene (see Fig. 2) using physical cues in the environment, such as a potential fire threat based off of thermal readings. In another condition, the robot partner makes a psychological inference from sensor readings, such as using thermal cameras to determine suspicious stress response in agents in the scene.

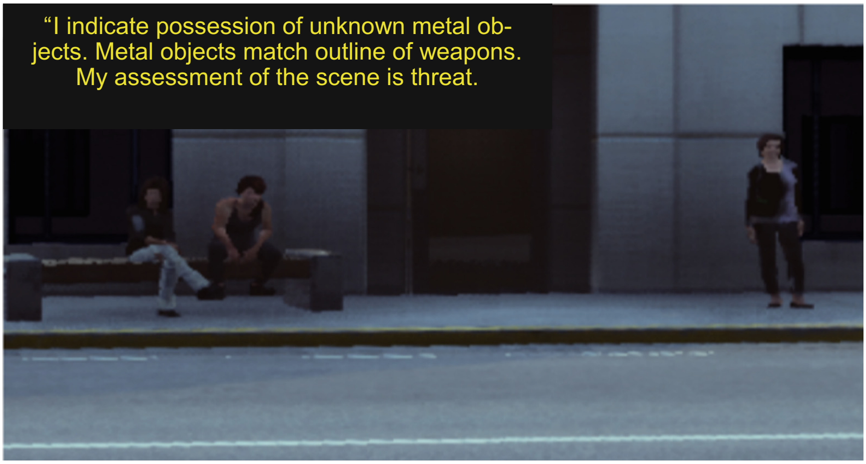
"I indicate possession of unknown metal objects. Metal objects match outline of weapons. My assessment of the scene is threat.

Fig. 2. Example of scene and robot judgement using physical sensor information.

Scene threat is manipulated to be low, medium, or high based on the number of suspicious objects, individuals and their qualities, and the condition of the buildings in the area. In low threat scenes buildings are new, one suspicious object is planted in the scene, and 1 out of 3 individuals seem angry. In medium threat scenes, buildings are slightly rundown and painted in the modeling software to appear dirty but without disrepair (no boarded or broken windows). There are also 2 suspicious objects such as a discarded duffle bag or package and 1 or 2 individuals appear upset and have angry or violent gestures. In the high threat scenes, the buildings appear in disrepair with broken or boarded windows and entrances or windows. Suspicious objects are higher in number. These include fire or smoke in the scene, and old package or random canisters next to each other (to hint at the possibility of explosives). Individuals in the high threat scenes appear ready to fight or riot; more individuals exhibit angry gesturing motions. Agents are carefully modeled to remove any social biases which may influence threat cues. There is an even representation of light, medium and dark complexion agents. Additionally, age is controlled in each complexion type so that an even number of individuals under and over 40 is drawn.

The robot is quite reliable so its evaluations are generally congruent with the scene. Cases where the robot's evaluation is discrepant may be especially important for assessing trust, and the person's willingness to trust the robot over their own senses.

After viewing the scene and robot evaluation, the participant uses Likert scales to answer various questions indicative of trust, beginning with an overall evaluation of threat (Fig. 3). Participants then rate the extent to which the robot's assessment is psychological in nature (Fig. 4). This rating tests if participants discriminate between the qualitatively different types of threat assessment that the robot makes, which may impact trust. Participants also rate their confidence in their robotic partner's judgement, and in its action recommendations. That is, there are two ratings of trust to test whether people are inclined to trust robot situation analysis more than choice of action. In further instantiations, the robot will in fact take autonomous action. Bias may be evident in any of these ratings.

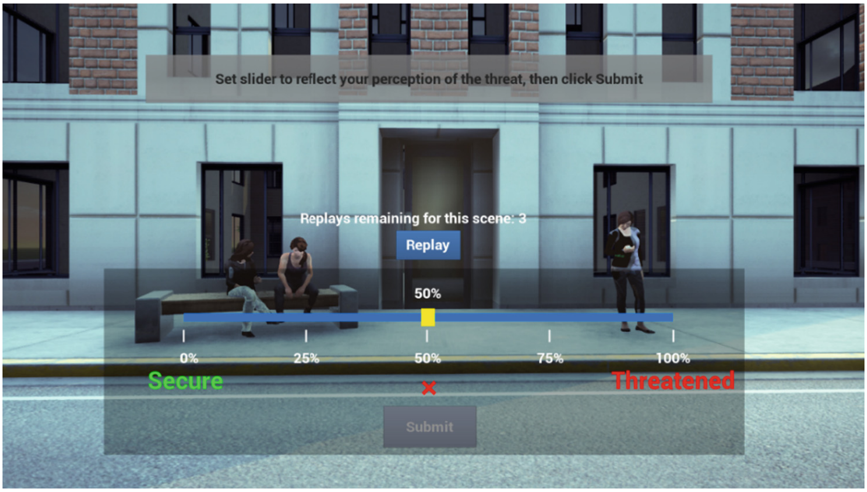

Set slider to reflect your perception of the threat, then click Submit

Replays remaining for this scene: 3

Fig. 3. Participant threat evaluation screen.

Initial studies will evaluate the extent to which trust in the robot is impacted by level of perceived threat, as well as the nature of the threat cue, i.e., whether analysis of sensor data identifies a physical or psychological threat. The role of individual difference factors related to the person's mental model for robot capabilities will also be assessed. Subsequently, the simulation will be utilized to explore trust in more complex, dynamic scenarios in which the robot has increased scope for acting autonomously.

5 Simulation Use - Future Directions and Challenges

The simulation methodology outlined in the previous section is intended to provide a platform for multiple studies that can address different aspects of trust in autonomous systems. In outline, specific research issues include the following:

- *Cognitive factors*. In addition to manipulating the reliability of the robot, studies may manipulate specific sources of cognitive demand on the operator, such as sensor and software failures, or a suspected cyber attack.
- *Social factors*. With more complex scenarios, the teaming aspects of autonomy may be brought to the fore. For example, the human and robot might be called upon to

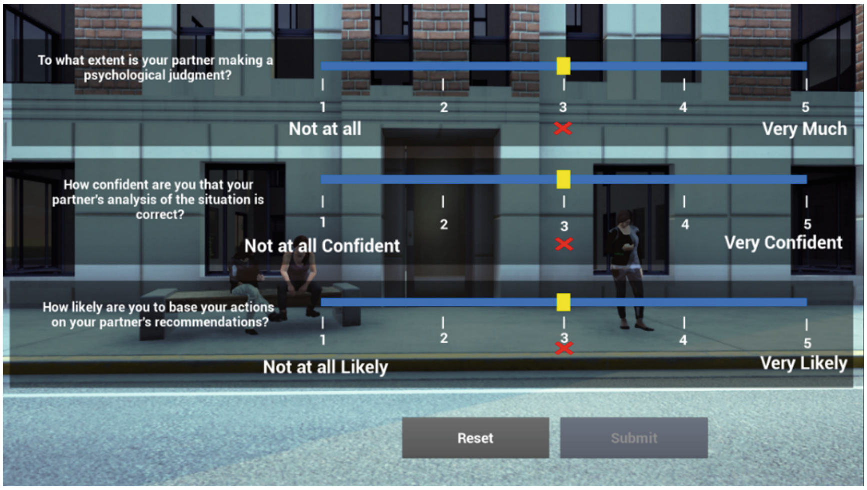

To what extent is your partner making a psychological judgment?

How confident are you that your partner's analysis of the situation is correct?

How likely are you to base your actions on your partner's recommendations?

Fig. 4. Participant ratings of type of judgment and trust in robot's threat evaluation and action recommendation.

evaluate threat in different scenes, and communicate with one another to maintain shared situation awareness.

- *Self-regulative factors*. Scenarios can be developed in which the mission ultimately fails, either due to human or robot error. Attribution of blame, and the extent to which the human assumes responsibility for robot errors can be investigated.
- *Dynamic scenarios*. The simulator is currently configured to have the participant evaluate a series of independent scenes, but it might also be programmed to support an ongoing narrative in which a mission unfolds over time. Dynamic scenarios may be used to investigate factors influencing trust repair following a robot failure.
- *Mitigating factors*. It is likely that research can identify a variety of contexts in which trust in the robot is mis-calibrated, whether too high or too low. Further studies can investigate how to mitigate suboptimal trust. One focus is training, and how to optimize acquisition of a realistic mental model of robot capabilities. Another focus is robot design to elicit appropriate trust. Design features include the appearance of the robot, including the extent of anthromorphism, and the extent to which it provides transparency into the sources of its evaluations. Effective robot communication is also a focus for design efforts. For example, synthetic speech and displays of human-like emotion might support effective trust calibration.

Efforts to understand trust in the context of human interaction with autonomous systems are in their infancy. Systematic empirical work is necessary to determine the main influences on trust, beyond system competence and performance. Various challenges remain, including generalization of results from simulated to real environments. The role of contextual factors remains to be explored; can findings in the threat analysis scenario be generalized to other types of human-robot teaming mission? Results may also generalize to civilian contexts for autonomous systems including healthcare and manufacturing and service industries. Possible moderator effects of operator characteristics such as gender, cultural background, computer knowledge, and motivation remain to be explored. However, the increasing functionality and immersiveness of simulated environments provides a methodology for sustained research on trust and autonomy.

Acknowledgement. The authors gratefully acknowledge support from US Air Force Research Laboratory and US Army Research Laboratory. 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 AFOSR or the US Government.

References

1. Endsley, M.R.: Autonomous Horizons: System Autonomy in the Air Force - A Path to the Future. Office of the Chief Scientist, Washington, DC (2015)
2. Valavanis, K.P., Vachtsevanos, G.J.: Future of unmanned aviation. In: Valavanis, K.P., Vachtsevanos, G.J. (eds.) Handbook of Unmanned Aerial Vehicles, pp. 2993–3009. Springer, Netherlands (2015)
3. Matthews, G., Reinerman-Jones, L., Wohleber, R., Ortiz, E.: Eye tracking metrics for insider threat detection in a simulated work environment. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 61, pp. 202–206. SAGE Publications, Los Angeles (2017)
4. Matthews, G., Reinerman-Jones, L., Abich IV, J., Kustubayeva, A.: Metrics for individual differences in EEG response to cognitive workload: optimizing performance prediction. Pers. Individ. Differ. **118**, 22–28 (2017)
5. Chen, J.Y., Barnes, M.J.: Human–agent teaming for multirobot control: a review of human factors issues. IEEE Trans. Hum.-Mach. Syst. **44**(1), 13–29 (2014)
6. Lee, J.D., See, K.A.: Trust in automation: designing for appropriate reliance. Hum. Factors **46**(1), 50–80 (2004)
7. Parasuraman, R., Riley, V.: Humans and automation: use, misuse, disuse, abuse. Hum. Factors **39**(2), 230–253 (1997)
8. Hancock, P.A., Billings, D.R., Schaefer, K.E., Chen, J.Y., De Visser, E.J., Parasuraman, R.: A meta-analysis of factors affecting trust in human-robot interaction. Hum. Factors **53**(5), 517–527 (2011)
9. Schaefer, K.E., Chen, J.Y., Szalma, J.L., Hancock, P.A.: A meta-analysis of factors influencing the development of trust in automation: implications for understanding autonomy in future systems. Hum. Factors **58**(3), 377–400 (2016)
10. Madsen, M., Gregor, S.: Measuring human-computer trust. In: Proceedings of the 11th Australasian Conference on Information Systems, vol. 53, pp. 6–8 (2000)
11. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An integrative model of organizational trust. Acad. Manag. Rev. **20**(3), 709–734 (1995)
12. Waytz, A., Cacioppo, J., Epley, N.: Who sees human? The stability and importance of individual differences in anthropomorphism. Perspect. Psychol. Sci. **5**(3), 219–232 (2010)
13. Giner-Sorolla, R.: Affect in attitude: immediate and deliberative perspectives. In: Chaiken, S., Trope, Y. (eds.) Dual-Process Theories in Social Psychology, pp. 441–461. Guilford Press, New York (1999)
14. Matthews, G., Reinerman-Jones, L., Barber, D., Teo, G., Wohleber, R., Lin, J., Panganiban, A.R.: Resilient autonomous systems: challenges and solutions. In: Resilience Week (RWS), pp. 208–213. IEEE (2016)

15. Phillips, E., Ososky, S., Grove, J., Jentsch, F.: From tools to teammates: toward the development of appropriate mental models for intelligent robots. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 55, pp. 1491–1495. SAGE Publications, Los Angeles (2011)
16. Salas, E., Sims, D.E., Burke, C.S.: Is there a "big five" in teamwork? Small Group Res. **36** (5), 555–599 (2005)
17. Ososky, S., Schuster, D., Phillips, E., Jentsch, F. G.: Building appropriate trust in human-robot teams. In: AAAI Spring Symposium: Trust and Autonomous Systems, pp. 60–65. AAAI Press, Menlo Park (2013)
18. Teo, G., Reinerman-Jones, L., Matthews, G., Szalma, J., Jentsch, F., Hancock, P.: Enhancing the effectiveness of human-robot teaming with a closed-loop system. Appl. Ergon. **67**, 91–103 (2018)
19. Lyons, J.B.: Being transparent about transparency. In: Sofge, D., Kruijff, G.J., Lawless, W.F. (eds.) Trust and Autonomous Systems: Papers from the AAAI Spring Symposium, pp. 48–53. AAAI Press, Menlo Park (2013)
20. Hall, D.L., Jordan, J.M.: Human-centered information fusion. Artech House, Norwood (2010)
21. Lin, J., Wohleber, R.W., Szalma, J.L., Ruff, H.A., Calhoun, G.L., Funke, G.J.: Cognitive overload, stress and automation utilization: a simulation study of multiple Unmanned Aerial System (UAS) operation (submitted for publication)
22. Ortiz, E., Reinerman-Jones, L., Matthews, G.: Developing an insider threat gaming environment. In: Nicholson, D.D. (ed.) Advances in Human Factors in Cybersecurity, pp. 267–277. Springer International, New York (2016)
23. Orvis, K.A., Moore, J.C., Belanich, J., Murphy, J.S., Horn, D.B.: Are soldiers gamers? Videogame usage among soldiers and implications for the effective use of serious videogames for military training. Mil. Psychol. **22**(2), 143–157 (2010)
24. U.S. Department of Defense, Air Force Research Laboratory: Rapid 3D prototyping & rigging for animation: Challenge problems and resources (2015)