



ARL-TR-7919 • JAN 2017



Humans and Autonomy: Implications of Shared Decision-Making for Military Operations

by Michael J Barnes, Jessie Y C Chen, and Susan Hill

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



Humans and Autonomy: Implications of Shared Decision-Making for Military Operations

by Michael J Barnes, Jessie Y C Chen, and Susan Hill
Human Research and Engineering Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) January 2017		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To) 1-31 October 2016	
4. TITLE AND SUBTITLE Humans and Autonomy: Implications of Shared Decision-Making for Military Operations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Michael J Barnes, Jessie Y C Chen, and Susan Hill				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Research Laboratory ATTN: RDRL-HRB-DE Aberdeen Proving Ground, MD 21005-5425				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-7919	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Issues related to defining the Soldier's role in future battlefields populated by autonomous systems are driving important US Army research programs. Mixed-initiative systems entailing shared decision-making between humans and intelligent software are a promising strategy that combines the advantages of human insight and autonomous control. This report discusses empirical results related to shared decision-making in the context of military applications including outcomes from research on intelligent agents, control of multiple unmanned systems, trust and transparency, cognitive architectures, natural language processing, and bi-directional interfaces. Overall, mixed-initiative systems show great promise, but more research will be required before such systems become part of large-scale operational environments. Effects of emotional response to autonomous systems, ethical software constraints, and machine learning transparency are identified as future research opportunities.					
15. SUBJECT TERMS autonomy, human factors, intelligent agents, controlled language processing, transparency, naturalistic interfaces					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
Unclassified	Unclassified	Unclassified	UU	42	Michael J Barnes 520-538-4702

Contents

List of Figures	iv
1. Introduction	1
1.1 Military Constraints	2
1.2 Mixed-Initiative Systems and a General Framework	3
1.3 The Conundrum of Control	5
2. RoboLeader and Human-Agent Control Processes	7
3. Trust and Transparency: Situation-Based Agent Transparency Model	8
3.1 Autonomy Research Pilot Initiative and Agent Transparency Research	9
3.2 Transparency and Trust: ARPI Results	10
4. Team Communications	11
4.1 ARL Robotic Collaborative Technology Alliance and Computational Cognitive Models	12
4.2 Language Processing	13
4.2.1 Command Processing	14
4.2.2 Controlled Processing	14
4.3 Graphic- and Video-Mediated Communications	15
4.4 Summary of Teaming Requirements	17
5. Naturalistic Interfaces	17
6. Summary and Discussion	19
7. Conclusions	22
8. References	23
List of Symbols, Abbreviations, and Acronyms	34
Distribution List	35

Approved for public release; distribution is unlimited.

List of Figures

Fig. 1	Characteristics of human-agent shared decision-making	4
Fig. 2	Control structures for human agent teams. Robots without tools are supervisor robots, while robots with tools at their base are operational robots.....	6
Fig. 3	RoboLeader interface for convoy operations.....	8
Fig. 4	Features of the SAT model of agent transparency	9
Fig. 5	Percentage correct data for levels of the SAT model in the second experiment.....	11
Fig. 6	Bi-directional communications between IDF commanders and a UAV using static (snapshots) and anchor (permanent landmarks) graphics	16
Fig. 7	Multiple robotic control devices used in field exercise	18
Fig. 8	Haptic forces to move a robotic swarm in the desired configuration around a nominal structure.....	19

1. Introduction

Combat requires systems that respond rapidly, efficiently, and safely while attaining mission objectives in situations that are increasingly complex (Barnes et al. 2014). “Smart” technologies are becoming ubiquitous in the modern world, changing the relationship of humans to their machines (*Economist*, 2016). In particular, autonomous systems are being developed for a wide variety of civilian and military applications to improve safety and reduce manpower restrictions (Purdy 2008; Greenemeier 2010; Murphy and Burke 2010; Osborn 2011; O’Dell 2013; Atherton 2015; Pellerin 2015; Brewster 2016). Before proceeding with the discussion of autonomous systems, it is important to note that using the term “autonomy” in military environments is possibly misleading. The dictionary definition of autonomy is “not subject to control from outside; independent, existing and functioning as an independent organism” (Dictionary.com 2016). Because of the uncertainties of the battlefield and the importance of human life, all combat systems are subject to ultimate human control (Chen and Barnes 2014; Endsley 2015b). Therefore, we define autonomous software agents in terms that are similar to the role that humans play in a military environment. The software agent is an intelligent nonhuman agent (IA) that has clear objectives, able to monitor its environment, and autonomous in the sense that it can generate courses of action (COAs) to obtain its objectives (Russell and Novrig 2009). However, as discussed in the following, the IA is always subordinate to its human supervisor much as a Soldier is subordinate to its commander.

Traditionally, humans and automated systems have been assigned to separate functions, but recent advances permit a fluid relationship approximating human teaming paradigms (Lyons and Stokes 2012; Cummings 2014). Various technologies are being designed to interact and communicate with human operators to ensure decision-making that is shared, flexible, and still human-centric (Fisher et al. 2007; Goodrich 2010; Goodrich et al. 2013; Chen and Barnes 2014). The differences in the following techniques involve the relative roles of humans and agents.

Adaptive systems monitor the environment or the operator’s cognitive state for triggering events. During overload or emergencies, control of system functions are automated. Similarly, the IA software relinquishes control to the operator during normal operations. Human control is ensured by an operator-initiated contract prior to the mission that defines decision precedence between the IA and the operator (Parasuraman et al. 2007; de Visser et al. 2011).

In contrast, **adjustable autonomy** (also referred to as adaptable autonomy) is automation that is instantiated at the discretion of the human during the mission. It

may take the form of “plays”, which are predetermined software solutions that the operator “calls” during the mission to address immediate tactical concerns (Miller and Parasuraman 2007).

In this report we will discuss **mixed-initiative systems**, in which the decision space is shared between IAs and human operators in real time. Mixed-initiative systems can incorporate both adaptive and adjustable software as part of the joint decision-making process (Hardin and Goodrich 2009, Goodrich 2010; Goodrich et al. 2013; Chen and Barnes 2014; Barnes et al. 2015).

To summarize, we are discussing mixed initiative autonomy wherein humans and IAs share decision-making. The IA has a degree of autonomy and communicates with its human supervisor. The focus of this report is on the decision-making relationship between IAs and its human supervisor(s) in future battle spaces that are dynamic, dangerous, and complex (Stone 2012). We will discuss human-agent teams in military environments as they relate to multiple systems, trust, transparency, agent architectures, 2-way communications, and the type of interfaces showing current progress and indicating possible future research opportunities (Stone 2012; Cummings 2014; Chen and Barnes 2014; Barnes et al. 2014; Endsley 2015a).

1.1 Military Constraints

The use of autonomy by the military raises special issues regarding rules of engagement (rules stipulating the circumstances under which use of weapons systems is permitted) beyond those involved in civilian applications (Defense Science Board 2012; Jentsch and Fincannon 2012). Many of these issues center on the trade-off between the utility of autonomy and its lethality (Singer 2010). As the development of technology affords new capabilities, there is ongoing concern that new autonomous capabilities may improve combat effectiveness at the risk of fratricide and civilian casualties (Barnes et al. 2014; Tiron 2003). The “fog of war” makes accidents inevitable, but at least initially the public will be far less forgiving if computer errors cause fatalities than if humans make the same mistakes. Mica Endsley (2015b), in her role as a Chief Scientist of the US Air Force, stresses that Department of Defense (DOD) Directive 3000.09 (2012) mandates safeguards for autonomous weapons, as shown in the following:

- “Semi-autonomous weapon systems that are onboard or integrated with unmanned platforms must be designed such that, in the event of degraded or lost communications, the system does not autonomously select and engage individual targets or specific target groups that have not been previously selected by an authorized human operator.”

- “The system design . . . addresses and minimizes the probability or consequences of failures that could lead to unintended engagements or to loss of control of the system.”
- “In order for operators to make informed and appropriate decisions in engaging targets, the interface between people and machines for autonomous and semi-autonomous weapon systems shall: (a) be readily understandable to trained operators, (b) provide traceable feedback on system status, and (c) provide clear procedures for trained operators to activate and deactivate system functions.”

These rules reinforce doctrine that military decisions are human responsibilities and operators must have a clear-cut understanding of the consequences of supervising autonomous systems. Furthermore, Endsley (2015b) concludes that autonomy must be integrated into the force structure. This entails training and the development of interfaces that make the autonomy understandable to ensure human-centered systems. Similarly, the Defense Science Board (DSB) (2012) emphasizes the importance of designing autonomy that fits into the network of military capabilities so as not to introduce “brittle” systems that have a negative impact on overall mission effectiveness. The DSB also stresses the importance of enablers of autonomy such as naturalistic interfaces to improve collaboration, trust, and situation awareness (SA) while reducing the Soldier’s physical and cognitive workload.

1.2 Mixed-Initiative Systems and a General Framework

Human-agent teaming is an important concept because it implies a personal relationship between the agent and the human. Ideally, the relationship will require bi-directional communications and a common worldview. IA architecture and its ability to communicate with humans is still primitive but it is progressing beyond the stage of literal translations and moving toward interpretation in terms of intent (Jurafsky and Martin 2009; Lomas et al. 2012; Wang et al. 2016). Figure 1 is a nominal human-agent framework intended to provide an overview of the issues that are discussed in the rest of the report.

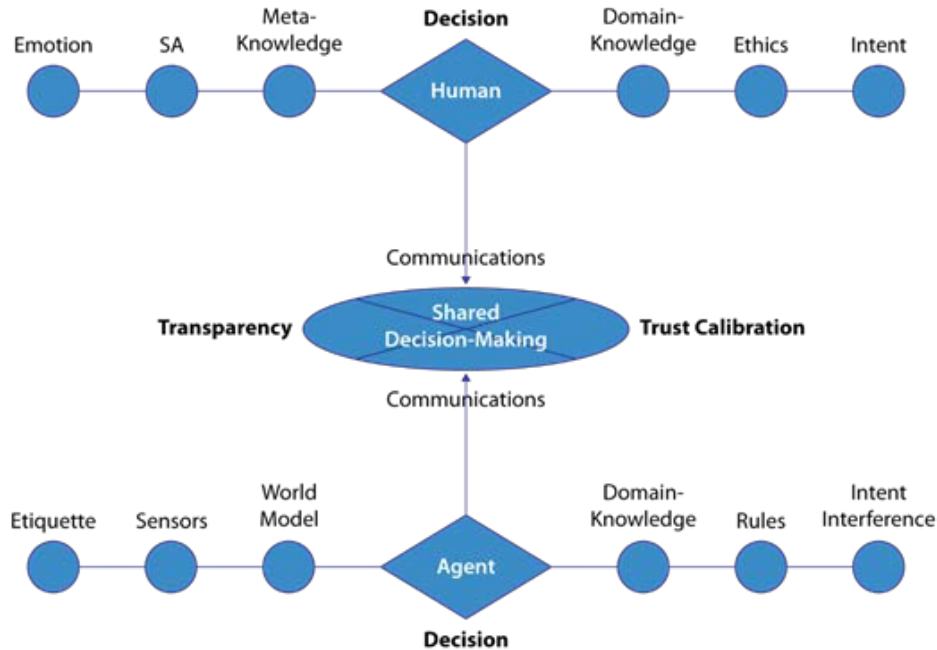


Fig. 1 Characteristics of human-agent shared decision-making

This framework will be used to motivate discussions of the various features influencing human-agent shared and separate decision spaces. The advantage of a human-agent partnership is that each element has its own strengths and weaknesses, and together they have the potential of being more effective than the sum of their parts (Chen and Barnes 2014). For example, the human will have greater meta-knowledge of political implications and changing strategic objectives, whereas the agent may have precise algorithms for specific technical challenges. On the other hand, there are problems in combining the elements into a cohesive decision structure. There are situations when human trust is misaligned with the agent’s reliability, causing humans to either over- or under-trust the agent’s decisions (Parasuraman and Riley 1997; Dzindolet et al. 2003; Lee and See 2004; Beck et al. 2007; Mercado et al. 2016). In similar fashion, the agent’s world model might be misaligned with the operator’s mental model and could misinterpret the intent of operator’s commands (Chen and Barnes 2014).

The user interface needs to be transparent so that agents and humans understand each other’s reasoning and uncertainties while making joint decisions (Lyons and Havig 2014). Creating mutual understanding requires calibrating the trust of the human operator and providing the IA with an ability to infer human intent (Mercado et al. 2016). There are distinct features of humans such as affect as well as cultural norms that must be accounted for in the agent’s behavioral repertoire (e.g., rules of etiquette) (Parasuraman and Miller 2004). To ensure safe operations, it is important to have protocols that address emergencies. For example, adaptive mechanisms

would permit the IA to react to dangerous situations such as collision avoidance without waiting for operator permission. Likewise, humans can take back authority for emergency situations (e.g., prevent fratricide) (Chen et al. 2010). It is crucial to define control procedures that are flexible but are consonant with operator's intent. Soldiers are encouraged to show initiative as long as they understand the intent of their original orders.

The remainder of this report discusses current research sponsored by the US Army Research Laboratory (ARL) (predominantly conducted in military environments) that addresses the following 5 issues: 1) control of multiple systems using intelligent agents, 2) requirements for developing transparency and appropriate trust necessary for human-agent interactions, 3) agent architectures and their implications for human-agent communication, 4) 2-way human-agent communications, and 5) naturalistic interfaces and their importance for efficient shared decision-making. In addition, we discuss areas of future research and the shortcomings of our current understanding of human-agent shared decision-making.

1.3 The Conundrum of Control

As mentioned previously, autonomy implies that the agent controls its own actions, but in a seeming contradiction we argue that ultimate control resides with the human operator. In the mixed initiative paradigm, it is necessary to develop protocols that dictate when the human, the agent, or both (collaborative) have decision precedence. The protocols can be mission-specific or be general in nature. However, in rapidly changing environments, human concurrence with agent decisions may not be practical. This is particularly true for multi-system control where the number of elements and the difficulties of controlling each element makes effective supervision difficult, if not impossible (Miller 1956; Lewis and Wang 2010; Schulte and Meitinger 2010, Chen et al. 2011; Lewis 2013). Metrics such as neglect time (time estimates of when supervisory attention is not needed for specific agents) and interaction time (average time that an operator needs to interact with an agent) are only useful if scheduling of attention by n-supervisors monitoring n-elements is predictable (Goodrich 2010; Goodrich et al. 2013). Combat by its nature is volatile and uncertain, making scheduling impractical in many situations.

A variety of strategies have proven effective in enhancing mixed-initiative decision-making. Some cognitive tasks are more amenable to automation than others. For example, information filtering/selection appears to be a good candidate for automation algorithms, but selecting an action with important consequences

usually requires human oversight (Parasuraman et al. 2000). There are also cases where autonomy can be assigned to various functions that are housekeeping (convoy separation) or because of time constraints (collision avoidance) (Wright et al. 2014). In other time-constrained situations, decision authority can be divided. For an incoming missile, an operator has the ability to override autonomy until a critical time limit is reached, after which missile defense systems kick in automatically (Parasuraman et al. 2007). In still other situations, such as identifying the importance of specific objects, the IA may detect an object but defer to the human to assess its significance (Jentsch and Fincannon 2012; Barnes et al. 2014). Autonomy needs to be flexible; authority resides with the operator but circumstances may require the IA to take the initiative.

The DSB (2012) suggested that autonomy will be particularly useful when an operator must interact with multiple assets. A caveat is that the number of assets controlled must be limited to that which can be managed effectively by a single human (Lewis and Wang 2010). To minimize control issues, ARL researchers are investigating an IA (Section 2) that acts as an intermediate supervisor by monitoring subordinate systems and by suggesting COA changes when unexpected events occur during the mission (Chen et al. 2011; Chen and Barnes 2012a,b; Chen and Barnes 2014). However, as Fig. 2 illustrates, hierarchical networks of agents can be expanded to enlist multiple local agents that interact with supervisory agents who in turn interact with the human operator. These paradigms use network technology with the human operator at the apex to reduce the problem space to manageable proportions without abrogating human decision authority (Hou et al. 2011; Chen and Barnes 2014). Multi-agent paradigms also have the advantage of being able to reconfigure the network as either the mission changes or an agent becomes disabled (DSB 2012).

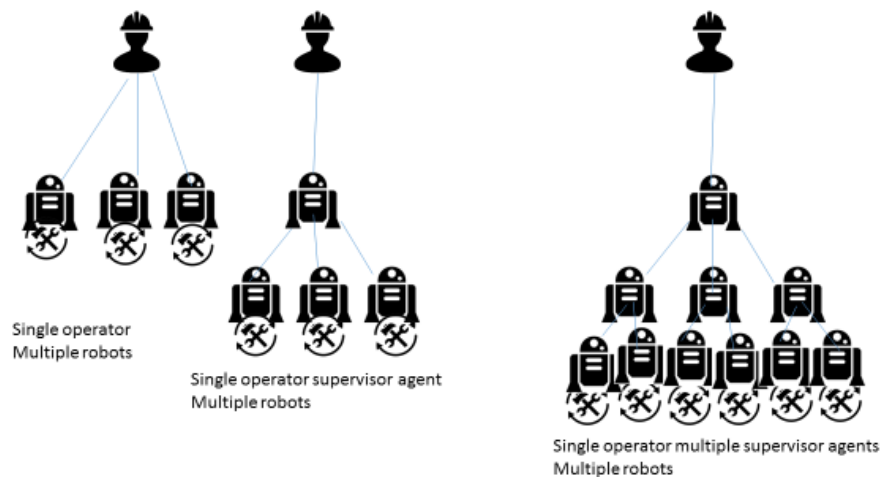


Fig. 2 Control structures for human agent teams. Robots without tools are supervisor robots, while robots with tools at their base are operational robots.

2. RoboLeader and Human-Agent Control Processes

RoboLeader is an ARL research paradigm investigating the human performance implications of instantiating an IA that supervises multiple systems and in turn is supervised by a human operator (Chen et al. 2010). RoboLeader researchers simulated various missions such as surveillance, policing, and convoy operations. Over the course of numerous studies, IA error type and rate, task difficulty, degree of autonomy, individual differences, agent transparency, and type of multitasking were investigated (Barnes et al. 2011; Chen and Barnes 2012a,b; Wright et al. 2014; Chen et al. 2016). The most recent studies simulated operators engaged in convoy operations in which they are supervising an unmanned ground vehicle (UGV), an unmanned aerial vehicle (UAV), and a manned ground vehicle while conducting 360° threat monitoring around their own vehicle (Wright et al. 2013, 2016 in press). The IA made convoy route change suggestions when unanticipated events occurred during the mission (Fig. 3). The results indicated the effects of varying reasoning information; succinct text explanations helped the operator reduce misuse of automation whereas supplying superfluous information hurt performance (Wright et al. 2016 in press). Multiple studies using the RoboLeader paradigm resulted in a better understanding of IAs contributions to military decision-making during manned/unmanned operations. Individual differences in gaming experience, spatial abilities, and an individual's confidence in attentional control proved to be ubiquitous factors in human-agent interactions, implying that training and decision support should be geared to individual aptitudes and experience rather than "one size fits all" solutions (Chen and Barnes 2011, 2014).

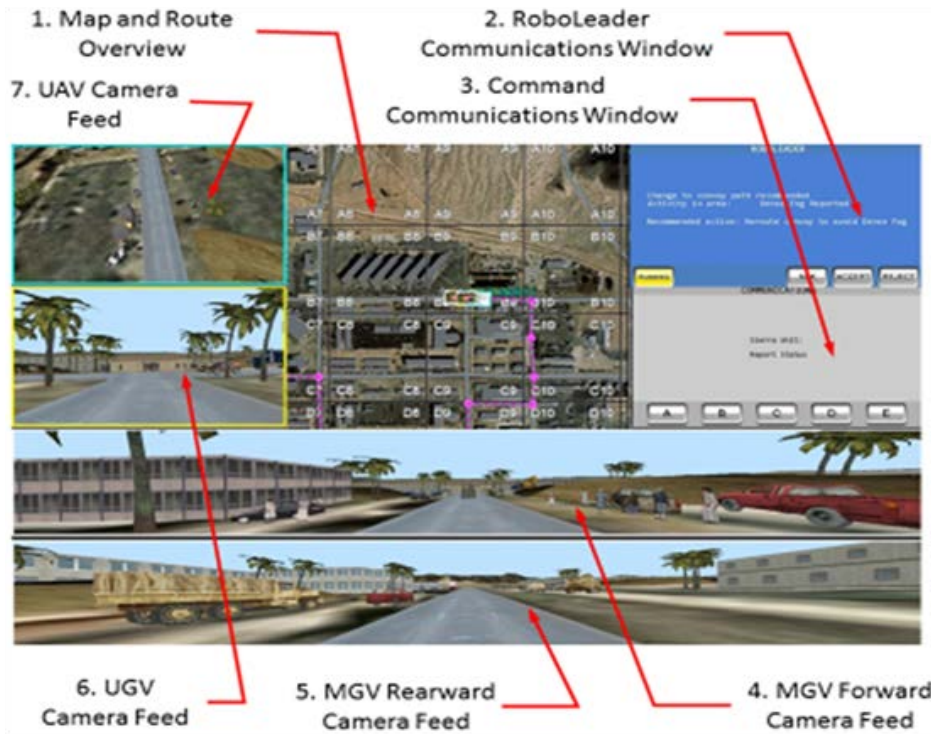


Fig. 3 RoboLeader interface for convoy operations (Wright et al. 2014)

3. Trust and Transparency: Situation-Based Agent Transparency Model

Trust is an important research topic for both automated and autonomous systems because it mediates between the reliability of such systems and operators' decisions to use them (Wickens 1994; Lyons and Havig 2014). Lee and See (2004) defined appropriate trust as human reliance on automation that minimizes disuse (failure to rely on reliable automation) and misuse (over-relying on unreliable automation) (Parasuraman and Riley 1997; Parasuraman and Manzey 2010; de Visser et al 2012). Trust can be measured either as an attitude (subjective measure) or as a behavior (misuse and disuse) (Lyons and Stokes 2012; Meyer and Lee 2013). Furthermore, trust can be a predisposition of the operator (trait) or depend on specific circumstances (state) (Schaeffer and Scribner 2015; Schaeffer et al. 2015). Subjective ratings have been shown to correlate with automation reliability, perception of the IA capabilities, and task difficulties as well as individual differences (Hancock et al. 2011; Schaeffer and Scribner 2015; Schaeffer et al. 2015). Lee (2012) suggests that in order to make the underpinnings of the automation algorithms transparent, the operator must be able to understand their purpose, process, and performance. Based on these and related concepts, ARL researchers (Chen et al. 2014; Chen and Barnes 2015) developed the SA-based

Agent Transparency (SAT) model (Fig. 4) to elucidate aspects of SA affecting trust (Endsley 1995, 2015b). SAT posits 3 transparency levels (Ls) of information to support the operator’s understanding of the IA’s decision process: L1) operator perception of the IA’s actions and plans, L2) comprehension of the IA’s reasoning process, and L3) understanding of the IA’s predicted outcomes. The purpose of the SAT model is to define the type of information necessary to give the operator insight into the IA’s intent, logic, and the perceived likelihood of obtaining its end state.

- To support operator’s **Level 1 SA** (*What’s going on and what is the agent trying to achieve?*)
 - *Purpose*
 - *Desire* (Goal selection)
 - *Process*
 - *Intentions* (Planning/Execution)
 - *Progress*
 - *Performance*
- To support operator’s **Level 2 SA** (*Why does the agent do it?*)
 - Reasoning process (*Belief*)(*Purpose*)
 - Environmental & other constraints
- To support operator’s **Level 3 SA** (*What should the operator expect to happen?*)
 - Projection to Future/End State
 - Potential limitations
 - Uncertainty; Likelihood of error
 - History of performance

Fig. 4 Features of the SAT model of agent transparency (Chen et al. 2014)

3.1 Autonomy Research Pilot Initiative and Agent Transparency Research

The US DOD Autonomy Research Pilot Initiative (ARPI) program is funded by the DOD to investigate the effects of autonomy in military environments and develop implementation practices based on acquired knowledge. The research is far-ranging and involves multiple programs. Two programs that ARL is involved in— Intelligent Multi-UxV Planner with Adaptive Collaborative/Control Technologies (IMPACT) and Autonomous Squad Member (ASM)—are discussed as exemplars of mixed initiative systems. Because of their realistic nature, time constraints, and complexity, the ARPI projects are ideal platforms for evaluating operator trust and, in particular, the efficacy of SAT model constructs. The IMPACT project is a tri-service program that investigates shared decision-making between multiple intelligent systems and the human operator for a base defense scenario in a littoral environment (Draper 2014). The operator’s role is to plan and supervise a mission with aerial, ground, and naval unmanned vehicles (UVs) that respond to suspicious activities related to base security (e.g., encroachments on the shore line). In typical mission scenarios, the base defense coordinator suggests an initial plan objective

referred to as a “play” (Miller and Parasuraman 2007). Then the IA/planner chooses available assets (UVs with sensor options) and generates the optimal route for the chosen platforms to attain the objective. During the mission, machine learning (ML) algorithms monitor progress. However, operators have final executive authority; they are able to tweak the plan or choose an option other than the agent’s preferred option. The objective of the ASM project is to investigate capabilities of a robotic asset to support squad-level performance during infantry missions (Chen and Barnes 2015; Selkowitz et al. 2016 in press). The ASM robot responds autonomously by adapting to squad tactics. Here again, Soldiers are the final arbitrators of the ASM decisions but are constrained by the fact that the ASM must respond rapidly to perturbations during the mission.

3.2 Transparency and Trust: ARPI Results

Mercado et al. (2016) tested implications of the SAT model simulating a simplified version of the IMPACT interface and missions. The experiment consisted of 24 simulated IMPACT scenarios (counterbalanced among participants) broken down into 8 scenarios for each SAT condition (L1, L1+2, and L1+2+3). Two COAs, A and B, were displayed for each scenario with option A being the IA’s preferred COA. During the experiment, alerts were given to participants that supported option A in 63% of the scenarios and B in 37%. Performance improved as a function of increasing SAT levels; misuse and disuse both decreased for the 2 higher transparency level conditions. Furthermore, subjective measures of trust increased for SAT Level 1+2+3, showing that attitude as well as performance were positively affected when the operator was given projected outcome information (Meyer and Lee 2013). Unlike previous research, neither workload nor response latency measures were degraded by the additional information comprising higher SAT levels (Helldin 2014). Of particular note was the finding that portraying uncertainty information in the L3 conditions improved the participants’ subjective trust (Mercado et al. 2016).

In a follow-on study, uncertainty (U) was parsed from L3 and the experimental conditions included L1+2 vs. L1+2+3 vs. L1+2+3+U (Stowers et al. 2016). Overall, uncertainty information improved operator performance. However, while the improvement in percentage correct was most evident in correct rejections (rejecting the suggested COA when the alternative was correct), Fig. 5 indicates that proper use (choosing the IA suggestion when it was correct) followed the same trend. Unlike the first experiment, measuring participants using subjective trust scales did not show significant improvements when uncertainty information was displayed to the operator (Chen et al. 2016; Stowers et al. 2016).




Fig. 5 Percentage correct data for levels of the SAT model in the second experiment (Stowers et al. 2016)

For a different task environment involving autonomous squad members (small robots), the results showed that adding uncertainty information (L1+2+3) did not improve subjective trust, whereas reasoning cues (L1+2) improved trust over baseline conditions (L1) (Chen and Barnes 2015). A follow-on experiment showed improvements in SA and trust when prediction information was available but showed no advantage to adding uncertain information (Selkowitz et al. 2016 in press). The experiments verified the main tenets of the SAT model. That is, the operator was better able to adjust the agent’s plan based on environmental or tactical changes because of the insights afforded by transparency information (Mercado et al. 2016). The efficacy of adding uncertainty to prediction was unclear especially in the ASM environment (Selkowitz et al. 2016 in press). This is at variance with other research that found that portraying uncertainty improved proper reliance on automated systems in a variety of environments. These findings indicate that more research is necessary to determine in which environments and for which display formats displaying uncertainty is beneficial (Bass et al. 2013; Bisantz 2013; Chen et al. 2014; Helldin 2014).

4. Team Communications

Besides understanding the IA’s decision processes, humans will need to interact with the agent as a team member to achieve effective shared decision-making (Green et al. 2008). Teams are defined as 2 or more entities that collaborate (share decision-making) and coordinate (synchronize tasks) to accomplish common goals. Especially in dynamic environments, an effective team requires compatible knowledge structures and the ability to communicate (Morrow and Fiore 2013). Both characteristics assume transparency among team members to the extent that

decisions and actions among team members are mutually understood. Knowledge structures can be shared mental models or specialized models by individual members, but the intent of individual team members must be communicated to other team members for effective joint action. In the case of teams of IAs and humans, the underlying knowledge structures may be quite different from those typically found in human-human teams. For example, in the IMPACT architecture, the IA processing is opaque to the operator but the resultant options can be graphically compared in terms of the trade-offs among the different outcomes (fuel consumption, time to target, etc.) indicating the IA's intended end state for each of the options (Behymer et al. 2015; Chen et al. 2016a,b).

Agent architectures are still a matter of considerable research interest (Chen and Barnes 2014). We discuss the trade-off between processing efficiency and transparency for shared problem-solving in subsequent sections. Unlike a human-human teaming relationship, natural dialogue is still difficult for unanticipated or unprogramed situations. However, progress in developing cognitive knowledge structures and natural language processing is making human-like interactions with agents increasingly likely for the near future (Lomas et al. 2012; *Economist* 2016).

4.1 ARL Robotic Collaborative Technology Alliance and Computational Cognitive Models

Cognitive models such as SOAR and ACT-R using rule-based systems, neural nets, and ML approaches were developed to model human information processing capabilities (Laird et al. 2011). Recently, cognitive models have been used to develop architectures to improve the robot's capacity to navigate in real-world environments and solve problems such as finding a doorway, locate a particular item in a room, or simulate an IA that acts as a surrogate crewmember (Ball et al. 2010; Kelley and McGhee 2013; Chen and Barnes 2014). There are a number of advantages to using cognitive approaches as opposed to purely algorithmic solutions. One is that cognitive models emulate a system that is adaptive and has proved successful in complex environments (i.e., the human brain), and another is that the similarities between the model's knowledge representations and human cognition should make transfer of information easier between agents and humans. ARL's Robotic Collaborative Technology Alliance has used a "find the backdoor" scenario both to develop the knowledge structures and the language processing required to control a small robot using simple commands to find a designated door.

Kelley and colleagues (e.g., Kelley and McGhee [2013]) have developed the Sub-Symbolic Robotics Intelligence System cognitive processing model to control robotic autonomous behaviors. Kelley and McGhee used the concept of episodic

memories to emulate cognitive processes such as using memory streams based on past experience (episodic memory) to address problems such as finding a back door in a building that required the robot to navigate around obstacles. The robotic IA builds software memories that are novel and those that are boring. The latter are memories that have little information because they do not change over time. The robotic agent combines the 2 types to remember when a boring memory transitions to a novel event in order to build a cognitive map to the correct back door solution. Other ARL research has used metaphors such as dreaming for knowledge consolidation and software constructs to represent emotions and temperament to make the robotic agent more accessible to its human teammate (Kelley 2014; Long et al. 2015). However, there is no reason to base agent intelligence solely on cognitive models; there are numerous useful agent technologies that are based on computational logic that solve specific problems efficiently (Fisher et al. 2007). Optimization algorithms such as Simultaneous Localization and Mapping have been used successfully for Army problems such as using robots to find the location of objects in buildings that would be unsafe for Soldiers performing the same function (Barnes et al. 2014). Whatever the IA's knowledge representation, it must translate to formats that humans understand to help foster a common language between the IA and its operator.

4.2 Language Processing

Although compatibility of knowledge structures is important for communication, compatibility by itself is not sufficient for 2-way communications. Language does not require either text or spoken dialogue; it does require syntax, semantics, and pragmatics to convey meaning during 2-way communications. Pragmatics assure that the communication is appropriate for the intended environment. For example, "go to the bank" has a different interpretation depending on whether the dialogue refers to natural surroundings or financial transactions (Jurafsky and Martin 2009). Thus natural dialogue is sensitive to nuance and intention and not simply its literal translation (Hoare and Parker 2010). This makes open-ended natural language processing (NLP) difficult and possibly impractical for some combat environments (Harris and Barber 2014). Chen and Barnes (2014) identified 3 gradations of language processing: command processing, controlled language processing, and open-ended NLP. The levels vary both in the size of their lexicons and the underlying sophistication of the software. As opposed to open-ended NLP, command and controlled language processing are both attuned to specialized tasking environments.

4.2.1 Command Processing

In practice, the command and the controlled processing strategies may overlap. However, controlled algorithms are geared to complex missions while command lexicons have been used successfully for limited task repertoires such as selecting menu options or controlling robotic movements (“move to the north of x”). The advantage of command processing is its simplicity and the ease in which its operators are able to assimilate its lexicon after limited training (Pettit et al. 2013; Barber et al. 2014). The lexicon is not limited to verbal utterances; successful interactions between humans and small robots have been demonstrated using gesture and tacton commands as well (Barber et al. 2013, 2015). Harris and Barber (2014) investigated various commercial off-the-shelf language processors to interpret speech for a limited domain lexicon for commanding small robots. The type of audio sensor and lexicon constraints influenced accuracy rate, which in general was fairly high (70%–80%). However, if the available online lexicon was much larger than the lexicon required to command the robot, they found the accuracy rate to be quite low (5%). The most likely reason was the increased likelihood of confusion between like-sounding utterances as the lexicon sized increased and became more open-ended. To standardize commands to the robot, recent efforts have focused on developing lexicons that represent Soldier speech patterns under realistic conditions (Barber et al. 2014), thus creating an easily learned lexicon of moderate dimensions.

4.2.2 Controlled Processing

More-sophisticated inference engines are needed for processors that enable 2-way communications that go beyond simple commands. Apple’s “Siri” and other commercial products indicate limited dialogues are possible with current technology. Also, because of their commercial potential, 2-way communication software will continue to improve, enabling human-agent interaction to mimic human-to-human dialogues in the near future. True dialogues involve intent inferencing and back and forth querying (such as, “Is this an object you wish me to investigate?” “No it is too oblong—check to the left about 10 meters.”) (Duplessis and Deviller 2015). Early agent architectures such as Belief-Desire-Intention modeled the agent’s processes in terms of beliefs (understanding of the environment), desires (objectives), and intentions (plan to achieve objectives) to capture an agent’s human-like qualities (Rao and Georgeff 1995; Chen and Barnes 2014). Two recent ARL-sponsored projects demonstrate progress toward more-mature language inferencing in military environments (Giammanco et al. 2015; Wang et al. 2016).

Controlled English (CE) is a specialized natural language representation developed by International Technology Alliance scientists from ARL and the United Kingdom (Giammanco et al. 2015; Xue et al. 2015). CE can be written by humans in English that is then translatable into a machine-readable format using domain semantics and predicate logic inferencing. Specifically, the user's conceptual model is written in CE as logical inference rules representing their relationships to specific military and civilian environments (Giammanco et al. 2015). For example, CE algorithms for military intelligence applications learn from interacting with intelligence analysts during real-world military vignettes. The interactions result in the ability of the CE to make sophisticated inferences concerning intelligence processes mimicking a human partner performing the same function (Mott et al. 2015). The crucial difference between CE and simpler command processes is the ability of CE algorithms to infer the meaning of environmental and situational cues.

Wang et al. (2016) simulated a self-explanatory agent that made its intentions and reasoning transparent to the operator using text-based messages during a simulated human-robot interaction task. The explanations were based on both stochastic reasoning by the agent and inferences about the human's preferences. Wang et al. described the algorithmic process as the following:

Decision-theoretic planning provides an agent with quantitative utility calculations that allow agents to assess trade-offs between alternative decisions under uncertainty. Recursive modeling gives the agents a theory of mind (Whiten 1991), allowing them to form beliefs about the human users' preferences, factor those preferences into the agent's own decisions, and update its beliefs in response to observations of the user's decisions.

Wang et al.'s agent not only provides information about its own reasoning process, but the IA attempts to understand the preference structure of its human teammate. Preliminary results suggest that participants were sensitive not only to the ability of the robot (its accuracy), but also to whether an agent generated an explanation for its action. Humans performed better (reduced misuse) and appropriately trusted even low-ability robots (reduced disuse) more often if participants understood the basis of the robot's decision. Future objectives include generating 2-way dialogues based on the inferencing and language processing abilities of the IA initially in simulation and eventually during exercises.

4.3 Graphic- and Video-Mediated Communications

In a field environment, communications using chat or voice may not be as efficient as graphic or video representations. Researchers at Ben-Gurion University (Oron-Gilad 2014) have collaborated with ARL to investigate the use of video feeds to

individual operators from both UGVs and UAVs surveilling possible insurgent activities for both stationary (a safe house) and mobile (car) targets. The results of the studies show distinct advantages to having video feeds from both ground and aerial views because of the advantages of giving the operator various perspectives of the ongoing mission (Oron-Gilad et al. 2011; Ophir-Arbelle et al. 2013; Barnes et al. 2014).

In follow-on research, Oron-Gilad and colleagues (2014) simulated bi-directional communication between a UAV and ground forces. A human played the role of the IA and human surrogates played the role of the military commanders directing a UAV. The purpose of the experimentation was to emulate graphic representations that can be updated by the UAV crew or its software and annotated by the ground commander (e.g., “Go to target X next”) during surveillance missions. In a recent simulation experiment (Oron-Gilad 2014), 9 participants with Israeli Defense Force (IDF) combat experience played the roles of the ground commanders interacting with a UAV operator. The UAV operator emulated a bi-directional agent capable of communicating with the mission commander by annotating imagery during the mission. Figure 6 shows 2 bi-directional graphics generating updates from the surrogate commanders and the UAV crewmember. The graphic on the left displayed important intelligence indicators that are annotated in real time, and the display refreshed itself as the mission progressed (series of static images). The graphic on the right showed images with anchor points indicating important map indicators as the mission progressed.

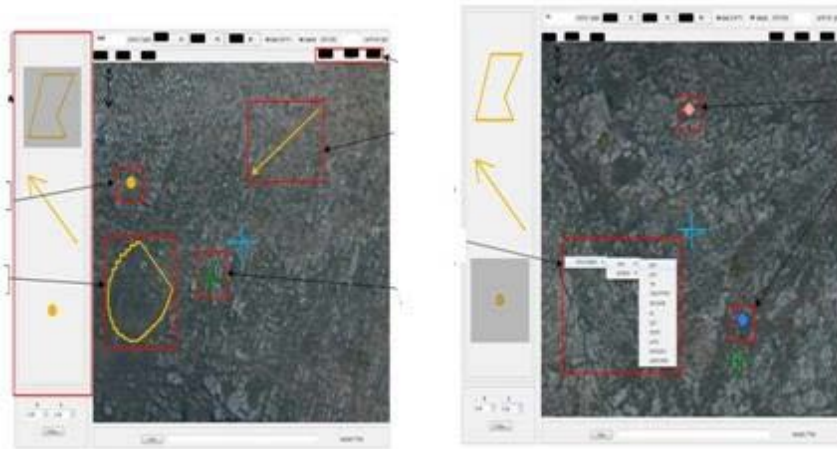


Fig. 6 Bi-directional communications between IDF commanders and a UAV using static (snapshots) and anchor (permanent landmarks) graphics

The participants found that combination of transitory snapshots and anchor points were viable sources of tactical information using bi-directional graphics as an interaction tool. Future research will investigate more complex missions using UGV as well as UAV videos. In subsequent experiments, enhanced bi-directional

interfaces using audio and tactile cueing will create a richer source of mission information. Eventually, the research plans include an IA with limited language processing abilities to replace the UAV operator in the bi-directional experiment. In a related effort (McDermott et al. 2015), the US Army is developing advanced visualization and analytics tools to autonomously search for and annotate videos with high-value intelligence. The system developers are in the process of integrating analytics such as face recognition tools and developing interfaces that allow the operator to query the system for 2-way interactions.

4.4 Summary of Teaming Requirements

Human-agent teams require humans to have insight into the IA's decision process and vice-versa (Chen and Barnes 2015). However, the IA must understand the implications of the human intentions that go beyond understanding the literal communications between team members (Hoare and Parker 2010). Bi-directional communications also require the give and take of normal conversations as each member of the team queries its teammate concerning the ongoing missions. The more closely the knowledge structures of the human-agent team are aligned, the more the necessity for extensive dialogue is ameliorated (Chen and Barnes 2014). This is important because military operations have additional constraints to minimize communications and to develop interfaces that are lightweight, quiet, easy to use, and non-observable (Barnes et al. 2014).

5. Naturalistic Interfaces

Shared decision-making interfaces will require advanced concepts to adhere to combat constraints especially for small-unit ground forces such as the ASM paradigm (Chen and Barnes 2015). A combination of multisensory interfaces improves the Soldier's ability to adjust to multiple situations such as the necessity for radio silence, day and night missions, and eyes-forward and hands-free. Hill et al. (2015) demonstrated the utility of controlling small robots using multiple control devices (stylus, voice, and glove) during an Army field experiment in 2014 (Fig. 7). The diversity of both input and display devices enabled communications with robots under a variety of field conditions.



Fig. 7 Multiple robotic control devices used in field exercise (Hill et al. 2015)

Elliott and her ARL colleagues (Elliott et al. 2010, 2015) have investigated voice, gesture, and tactile communications to develop naturalistic interfaces that enhance the ability of the operator to control robots. The advantages of these interfaces are that Soldiers are not burdened by having to look down at a display, hands are freed to carry weapons, and Soldiers can be signaled covertly concerning both their own and the robot's location. In a field experiment at Fort Benning, Georgia, Soldiers using a tactile vest were able to respond more rapidly and accurately than using traditional signaling methods during a reconnaissance mission (Elliott et al. 2016). For night missions in particular, tactile vests proved an important navigation aid (Pomranky-Hartnett et al. 2015). These findings reinforce the corpus of tactile research indicating the robustness of tactile communication under a variety of conditions (Van Erp 2007; Jones and Sarter 2008; Elliott et al. 2009; Barnes et al. 2014; Barber et al. 2015; Elliott et al. 2015).

The technology for gesture control systems is undergoing rapid development with organizations exploring multiple options because of the perceived commercial benefit of these devices (Elliott et al. 2016). The 2 principal approaches are camera-based and instrumented gloves (wireless). Both approaches assume the operator is able to signal unambiguously, and both types depend on algorithms embedded in the asset (e.g., Hidden Markov Models) to disambiguate signals. The results related to gesture control are mixed. For example, Soldiers in Elliott et al.'s (2015) experiment rated the utility of instrumented gloves highly but were more accurate using the tablet display for robot control. Currently, various problems with gesture control have been noted. The systems are too bulky to be practical, signaling by the operator can be difficult, and there are security issues related to wireless transmissions (Elliott et al. 2016). Gesture control's most likely use will be in conjunction with voice, visual, and tactile interfaces to supplement the difficulties

and advantages of each modality for particular situations. The caveat is that the packaging of multimodal solutions will have to become feasible (durable, wearable, and lightweight) for military applications (Hill et al. 2015).

The advantage to multisensory interfaces is that they offer both practical and performance enhancing methods for communication between IAs (e.g., robots) and humans during complex military operations (Hancock et al. 2011; Barnes et al. 2014). Nontraditional interfaces are being investigated to enable human control of new technologies such as robotic swarms that are holistic configurations of many agents that reconfigure autonomously. In a Brigham Young University study (Alder et al. 2015), a swarm is being developed that is controlled by a haptic interface that uses pressure and directional cues to “drag” the swarm around objects representing buildings. The user interface is designed to give the operator intuitive feedback during the drag operation that does not require “heads-up” visual cues. Figure 8 shows the haptic interface control dynamics necessary to move the swarm around a simulated structure (Alder et al. 2015).




Fig. 8 Haptic forces to move a robotic swarm in the desired configuration around a nominal structure (Alder et al. 2015)

6. Summary and Discussion

We reviewed human-agent research focusing on shared decision-making in which humans supervise multiple IAs that have varying degrees of autonomy (Chen et al. 2011; Draper 2014). The scope of the report encompasses the following 5 areas of mixed initiative decision-making and their enabling technologies:

Approved for public release; distribution is unlimited.

1. Control issues related to using IAs as intermediate supervisors to monitor n-systems that, in turn, are supervised by a human operator (Chen et al. 2011; Chen and Barnes 2012a; Chen and Barnes 2014).
2. Transparency and trust issues related to the operator's understanding of the IA's intent, reasoning, and projected outcomes (Chen et al. 2014; Mercado et al. 2016; Selkowitz et al. 2016 in press; Stowers et al. 2016).
3. Shared knowledge structures including computational cognitive architectures that enable effective collaboration (Kelley 2014).
4. Language processing software to foster 2-way communications between agents and humans (Giammanco et al. 2015; Wang et al. 2016).
5. Specialized interfaces to expedite control of embodied agents using gesture, voice, and haptic controllers (Pettit et al. 2013; Barnes et al. 2014; Alder et al. 2015; Elliott et al. 2015, 2016).

Progress in integrating these components into systems that approximate human teams is encouraging, but the state-of-art is still very much in the research phase (Chen and Barnes 2014). Needless to say, there are important issues of human-agent collaboration that go beyond the research covered in this report. We briefly review some of these areas to discuss pertinent issues as grist for future research.

Eventually, autonomous robots will be used for multiple functions including security, medical uses, maintenance, and transportation as well as combat roles that will require IAs to be part of the Soldier's daily experience. Especially for embodied agents such as ground robots, IAs must fit into a social network that requires them to interact within human moral and emotional expectations (Jones and Schmidlin 2011; Fiore et al. 2013). Humans and robots will have to co-exist and respect each other's space, which will require developing an etiquette to guide their interaction (Parasuraman and Miller 2004). IAs that collaborate with operators on a personal level should be designed to express and respond to nonlinguistic cues. Breazeal (2003, 2009) discussed the benefits and challenges of designing an anthropomorphic robot ("Kismet") whose facial expression is able to convey emotional cues to facilitate a more natural relationship between the robot and its human clients. Poorly designed IAs can have negative effects as well, causing distrust and reluctance to share the same living space (Breazeal 2009).

Arkin and Ulam (2009) believe that ethical considerations need to be encoded into robotic architectures to ensure that autonomy does not lead to dangerous behaviors. During military operations, autonomous decisions need to be made rapidly and inflexible rules could be disastrous. However, as long as software rules are transparent, having ethical brakes embedded into IA architectures will give its

human supervisor more oversight during dangerous situations (Barnes and Evans 2010). Scheutz (2016 in press) argues that designing robots that avoid harming humans is not sufficient (i.e., “implicitly moral” robot). Robots need to be “explicitly moral”, using the same reasoning process humans make when encountering a difficult moral dilemma. Obviously, this is not always easy for either a human or an IA, but a robot can signal its operator that the actions it is being directed to perform have ethical consequences. For example, Briggs and Scheutz (2012) investigated human responses to a robot showing discomfort concerning destroying simulated buildings. The results indicate that robots can influence their human supervisors’ ethical decisions through nonexplicit behaviors. Whether the simulated advisories would be effective during real-world operations is unclear. However, an IA advising against executing an unethical or dangerous option would remind the operator of the consequences of the proposed actions (possibly via displaying projected outcomes as suggested in the SAT model–Level 3 [Chen et al. 2014]). Advisories will be particularly effective if it is clear to the operator that the advisories are based on military doctrine (Endsley 2015b). The trade-off between speed and control, lethality and safety, and ethics and expediency are important issues that will permeate future IA research (*Economist* 2016).

An important trend in agent technology is the greater use of ML methods. ML is not a single approach but rather subsumes many algorithmic and statistical methods such as evolutionary algorithms, Bayesian approaches, adaptive control theory, neural nets, and the like. Nilsson (1998), in his now classical textbook, defined ML as “... a machine learns whenever it changes its structure, program, or data (based on its inputs or in response to external information) in such a manner that its expected future performance improves”. There are 2 principal issues related to ML: 1) the underlying algorithmic approach is often opaque and 2) the reasoning for choosing COA will change as more information is accumulated. The underlying logic depends on the technique used. For example, evolutionary algorithms use a fitness function (e.g. number of casualties) to choose a solution at each iteration (Suantak et al. 2001). In this case, the efficacy of the ML solution can be described in terms of its expected outcomes (in relation to its fitness function) as the COA changes over time. Because of the dynamics of military environments, ML will be an essential tool for IA technology (Draper 2014). Both IAs and their human counterparts must adapt to military environments that are in a state of continual flux. However, it is difficult to believe that humans will have trust in a system that changes its preferred solutions unless the logic and expected outcomes driving the changes are transparent (Chen et al. 2014).

7. Conclusions

We conclude that shared decision-making between humans and IAs shows potential as an effective means of addressing the complexities of modern warfare. However, more research is needed before this potential is realized. The human's and the agent's perception of the world must be aligned to create the synergy necessary to take full advantage of their respective strengths and limitations. This requires transparency of both the human and the agent's intent, logic, and projected outcomes. The agent's software architecture must support both bi-directional transparency and human-agent communication. Language processing from simple commands to complex inferencing is maturing rapidly, making human-agent teams that can communicate with each other feasible in the near future. Future research efforts should address the effects of emotions on human-agent team building, ethical constraints of autonomy, and the promise and perils of machine learning.

8. References

- Alder CK, McDonald SJ, Colton MB, and Goodrich MA. Toward haptic-based management of small swarms in cordon and patrol. Presented at the Swarm/Human Blended Intelligent Workshop; 2015 Sep; Cleveland, OH.
- Arkin RC, Ulam P. An ethical adapter: behavioral modification derived from moral emotions. Presented at the IEEE International Symposium on Computational Intelligence in Robotics and Automation; 2009 Dec 15–18; Daejeon, South Korea.
- Atherton KD. Echo seeker submarine patrols the seafloor on its own. *Popular Science*. 2015 Aug 10.
- Ball J, Myers C, Heiberg A, Cooke N, Matessa M, Freiman M, Rodgers S. The synthetic teammate project. *Comput Math Org Theory*. 2010;16:271–299.
- Barber D, Lackey S, Reinerman-Jones L, Hudson I. Visual and tactile interfaces for bi-directional human robot communication. *SPIE Defense, Security, and Sensing*. 2013:87410U.
- Barber D, Wohleber RW, Parchment A, Jentsch F, Elliott L. Development of a squad level vocabulary for human-robot interaction. In: Shumaker R, Lackey S, editors. *Virtual, augmented and mixed reality: designing and developing virtual and augmented environments*; 2014; vol. 8525; p. 139–148 [accessed 2016 Feb]. http://link.springer.com/10.1007/978-3-319-07458-m0_14.
- Barber D, Reinerman-Jones L, Matthews G. Learning and performance toward a tactile language for human-robot interaction: two studies of tacton. *Human Factors*. 2015;57(3):471–490.
- Barnes MJ, Evans AW. Soldier-robot teams in future battlefields: an overview. In: Barnes M, Jentsch F, editors. *Human-robot interactions in future military operations*. Farnham (UK): Ashgate; 2010. p. 9–29.
- Barnes MJ, Chen JYC, Jentsch F, Redden ES. Designing effective Soldier-robot teams in complex environments: training, interfaces, and individual differences. In: Harris D, editor. *Engineering psychology and cognitive ergonomics: lecture notes in computer science*. Vol. 6781. Berlin (Germany): Springer-Verlag; 2011. p. 484–493.
- Barnes MJ, Chen JYC, Jenstch F, Oron-Gilad T, Redden E, Elliott L, Evans A. *Designing for humans in autonomous systems: military applications*. Aberdeen

- Proving Ground (MD): Army Research Laboratory (US); 2014 Jan. Report No.: ARL-TR-6782.
- Barnes MJ, Chen JYC, Jentsch F. Designing for mixed-initiative interactions between humans and autonomous systems in complex environments. Proceedings of 2015 IEEE Conference on Systems, Man, and Cybernetics (SMC); 2015 Oct; Hong Kong. p. 1386–1390.
- Bass EJ, Baumgart LA, Shepley KK. The effect of information analysis automation display content on human judgment performance in noisy environments. *Journal of Cognitive Engineering and Decision-making*. 2013;7(1):49–65.
- Beck HP, Dzindolet MT, Pierce LG. Automation usage decisions: controlling intent and appraisal errors in a target detection task. *Human Factors*. 2007;49:429–437.
- Behymer KJ, Mersch EM, Ruff HA, Calhoun GL, Spriggs SE. Unmanned vehicle plan comparison visualization for effective human-autonomy teaming. Presented at the 6th Applied Human Factors and Ergonomics International Conference, 2015 July 26–30, Las Vegas, NV.
- Bisantz A. Uncertainty visualization and related topics. In: Lee J, Kirlik A, editors. *The Oxford handbook of cognitive engineering*. Oxford (UK): Oxford University Press; 2013.
- Breazeal C. Emotion and sociable robots. *Int J Hum-Comp Stud*. 2003;59:119–155.
- Breazeal C. Role of expressive behavior for robots that learn from people. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2009;364(1535):3537–3538.
- Brewster S. The age of the autonomous robot is upon us. *Fortune*; 2016 Mar 29 [accessed 2016 Apr 9]. www.fortune.com/2016/03/29/autonomous-robots-startups.
- Briggs G, Scheutz M. Investigating the effects of robotic display of protest and distress. Proceedings of the 2012 Conference on Social Robotics; 2012 Oct; Chengdu, China.
- Chen JYC, Barnes MJ, Qu Z. RoboLeader: a surrogate for enhancing the human control of a team of robots. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2010 Jan. Report No.: ARL-MR-0735.
- Chen JC, Barnes MJ, Harper-Sciarini M. Supervisory control of multiple robots: human performance issues and user interface design. *IEEE Transactions on*

- Systems, Man, and Cybernetics – Part C: Applications and Reviews. 2011;41(4):435–454.
- Chen, JYC, Barnes MJ. Supervisory control of multiple robots in dynamic tasking environments. *Ergonomics*. 2012a;55:1043–1058.
- Chen, JYC, Barnes MJ. Supervisory control of multiple robots: effects of imperfect automation and individual differences. *Human Factors*. 2012b;54(2):157–174.
- Chen JYC, Barnes MJ. Human–agent teaming for multirobot control: a review of human factors issues. *IEEE Transactions on Human-Machine Systems*. 2014;44(1):13–29.
- Chen JYC, Barnes MJ. Agent transparency for human-agent teaming effectiveness. *Proceedings of the IEEE conference on Systems, Man, and Cybernetics (SMC)*; 2015 Oct; Hong Kong.
- Chen JYC, Procci K, Boyce M, Wright J, Garcia A, Barnes MJ. Situation awareness-based agent transparency. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2014 Apr. Report No.: ARL-TR-6905.
- Chen JYC, Barnes MJ, Selkowitz A, Stowers K. Effects of agent transparency on human-autonomy teaming effectiveness. Presented at the 2016 IEEE Conference on Systems, Man, and Cybernetics (SMC); 2016a Oct 6–12, Budapest, Hungary.
- Chen JYC, Barnes M, Selkowitz A, Stowers K, Lakhmani S, Kasdaglis, N. Human-autonomy teaming and agent transparency. *Proceedings of the ACM Intelligent User Interface (IUI) Conference*; 2016b Mar 7–10; Sonoma, CA. p. 28–31.
- Cummings M. Man vs. machine or man + machine? Expert opinion: IEEE intelligent systems. *IEEE Computer Society*; 2014 Sep–Oct.
- Defense Science Board (DSB). Role of autonomy in DOD systems. Washington (DC): Office of the Undersecretary of Defense (US); 2012.
- de Visser E, Parasuraman R. Adaptive aiding of human-robot teaming: effects of imperfect automation on performance, trust and workload. *Journal of Cognitive Engineering and Decision-making*. 2011;5(2):2019–2023.
- de Visser E, Krueger F, McKnight P, Scheid S, Smith M, Chalk S, Parasuraman R. The world is not enough: trust in cognitive agents. *Proceedings of the 56th Annual Human Factors and Ergonomics Society (HFES) Meeting*; 2012 Oct 22–26; Boston, MA.

- Department of Defense. *Autonomy in weapon systems*. Washington (DC): Department of Defense (US); 2012 Nov 12.
- Draper M. *Realizing autonomy via intelligent adaptive hybrid control: adaptable autonomy for achieving UxV RSTA team decision superiority – year 1 report*. Dayton (OH): Air Force Research Laboratory (US); 2014.
- Dictionary.com. *Autonomous*. Oakland (CA) Dictionary.com LLC; 2016 [accessed Dec 2016]. <http://www.dictionary.com/browse/autonomous?s=t>.
- Dzindolet M, Peterson S, Pomranky R, Pierce L, Beck H. The role of trust in automation reliance. *Int J Hum-Comp Stud*. 2003;58:697–718.
- March of the machines: a special report on artificial intelligence. *The Economist*. 2016 July 1.
- Elliott L, Covert M, Redden E. Overview of meta-analyses investigating vibrotactile versus visual display options. In: Jacko JA, editor. *Human computer interaction: part II*. Berlin (Germany): Springer-Verlag; 2009. p. 435–441.
- Elliott L, van Erp J, Redden E, Duistermaat M. Field-based validation of a tactile navigation device. *IEEE Transactions on Haptics*. 2010;3(2):78–87.
- Elliott C, Jensen E, Redden S, Pettit E. *Robotic telepresence: perception, performance and user experience*. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2012 Feb. Report No.: ARL-TR-5928.
- Elliott L, Mortimer B, Hartnett-Pomranky G, Zets G, Mort G. *Augmenting Soldier situation awareness and navigation through tactile cueing*. Proceedings of the 17th International Conference on Human-Computer Interaction; 2015 Aug 2–7; Los Angeles, CA.
- Elliott L, Hill S, Barnes M. *Gesture-based controls for robots: overview and implications for use by Soldiers*. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2016 Jul. Report No.: ARL-TR-7715.
- Endsley MR. Toward a theory of situation awareness in dynamic systems. *Human Factors*. 1995;37:32–64.
- Endsley MR. Situation awareness misconceptions and misunderstandings. *Journal of Cognitive Engineering and Decision-making*. 2015a;9:4–15.
- Endsley MR. *Autonomous horizons: system autonomy in the air force – a path to the future, vol. I: human autonomy teaming*. Washington (DC): Department of the Air Force (US); 2015b.

- Fisher M, Bordini R, Hirsch H, Torroni B. Computational logic and agents. *Computational Intelligence*. 2007;23(1):61–91.
- Fiore SM, Wiltshire TJ, Lobato EJC, Jentsch F, Huang WH, Axelrod B. Towards understanding social cues and signals in human-robot interaction: effects of robot gaze and proxemic behavior. *Frontiers in Psychology: Cognitive Science*. 2013;4:859.
- Giammanco C, Mott M, McGowan R. Controlled English for critical thinking about the civil-military domain. Presented at the International Technology Alliance in Network and Information Sciences Annual Fall Meeting; 2015 Sep.
- Goodrich MA. On maximizing fan-out: towards controlling multiple unmanned vehicles. In: Barnes M, Jentsch F, editors. *Human-robot interactions in future military operations*. Farnham (UK): Ashgate; 2011. p. 375–395.
- Goodrich MA. Multitasking and multi-robot management. In: Lee JD, Kirlik A, editors. *Oxford handbook of cognitive engineering*. New York (NY): Oxford; 2013.
- Green SA, Billingham M, Chen X, Chase JG. Human-robot collaboration: a literature review and augmented reality approach in design. *International Journal of Advanced Robotic Systems*. 2008;5(1):1–18.
- Greenemeier L. Robot pack mule to carry loads for G.I.s on the move. *New York (NY): Scientific American*; 2010 [accessed 2015 Dec 15]. <http://www.scientificamerican.com/article.cfm?id=boston-dynamics-ls3>.
- Hancock PA, Billings DR, Schaefer KE, Chen JYC, de Visser E, Parasuraman R. A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*. 2011;53(5):517–527.
- Hardin B, Goodrich MA. On using mixed-initiative control: a perspective for managing large-scale robotic teams. *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction*; 2009 Mar 11–13; La Jolla, CA. p. 165–172.
- Harris J, Barber D. Speech and gesture interfaces for squad level human robot teaming. *Proceedings of the SPIE Defense, Security, and Sensing: Unmanned Systems Technology Conference*; 2014 May 5–9; Baltimore, MD.
- Helldin T. Transparency for future semi-automated systems [doctoral dissertation]. [Orebro (Sweden)]: Orebro University; 2014.

- Hill SG, Barber D, Evans AW. Achieving the vision of effective Soldier-teams: recent work in multimodal communications. Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction; 2015 Mar; Portland, OR.
- Hoare J, Parker L. Using on-line conditional random fields to determine human intent for peer-to-peer human robot training. Proceedings of the IEEE International Conference on Intelligent Robots and Systems; 2010 Oct 18–22; Taipei, Taiwan.
- Hou M, Zhou H, Zhou M, Arrabito G. Optimizing operator-agent interaction in intelligent adaptive interface design: a conceptual framework. IEEE Transactions on Systems, Man, and Cybernetics – Part C: Applications and Reviews. 2011;41(2):161–178.
- Jentsch F, Fincannon T. Multiple robot-multiple operator control and teamwork. Proceedings of the Information Technology Conference of the American Institute for Aeronautics and Astronautics; 2012; Orange County, CA.
- Jones KS, Schmidlin EA. Human-robot interaction: toward usable personal service robots. Reviews of Human Factors and Ergonomics. 2011;7:100–148.
- Jones LA, Sarter NB. Tactile displays: guidance for their design and application. Human Factors. 2008;50(1):90–111.
- Jurafsky D, Martin JH. Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition. Upper Saddle River (NJ): Prentice-Hall; 2009.
- Kelley TD. Robotic dreams: a computational justification for the post hoc processing of episodic memories. International Journal of Machine Consciousness. 2014;6(2):109–123.
- Kelley TD, McGhee S. Combining metric episodes with semantic event concepts within the symbolic and sub-symbolic robotics intelligence control system (SS-RICS). SPIE Defense, Security, and Sensing. 2013 May. p. 87560L.
- Laird J, Derbinsky N, Voigt J. Performance evaluation of declarative memory systems in soar. Proceedings of the Behavior Representation in Modeling and Simulation Conference; 2011 Mar 21–24; Sundance, UT. p. 33–40.
- Lee JD. Trust, trustworthiness, and trustability. Presentation at the Workshop on Human Machine Trust for Robust Autonomous Systems; 2012; Ocala, Florida.
- Lee JD, See KA. Trust in automation: designing for appropriate reliance. Human Factors. 2004;46(1):50–80.

- Lewis M. Human interaction with remote robots. *Human Factors and Ergonomics*. 2013;9:31–173.
- Lewis M, Wang J. Coordination and automation for controlling robot teams. In: Barnes MJ, Jentsch F, editors. *Human robot interactions in future military operations*. Farnham (UK): Ashgate; 2011. p. 398–418.
- Lomas M, Chevalier R, Cross E, Garrett C, Hoare J, Kopack M. Explaining robot behavior. *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*; 2012 Mar 5–8; Boston, MA.
- Long L, Kelley T, Avery E. An emotion and temperament model for cognitive mobile robot. Presented at the 24th Conference on Behavior Representation in Modeling and Simulation (BRIMS); 2015 Mar 31– Apr 3; Washington, DC.
- Lyons JB, Stokes CK. Human-human reliance in the context of automation. *Human Factors*. 2012;54(1):112–121.
- Lyons JB, Havig PR. Transparency in a human-machine context: approaches for fostering shared awareness/intent. In: Shumaker R, Lackey S, editors. *Virtual, augmented and mixed reality: designing and developing virtual and augmented environments*. Berlin (Germany): Springer; 2014. p. 181–190.
- McDermott P, Plott B, Ries A, Touryan J, Barnes M, Schweitzer K. Advanced video activity analytics (AVAA): human factors evaluation. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015 May. Report No.: ARL-TR-7286.
- Meyer and Lee J. Trust, reliance, compliance. In: Lee J, Kirlik A, editors. *The Oxford handbook of cognitive engineering*. Oxford (UK): Oxford University Press; 2013.
- Mercado J, Rupp M, Chen J, Barber D, Procci K, Barnes M. Intelligent agent transparency in human-agent teaming for multi-UxV management. *Human Factors*. 2016;58(3):401–415.
- Miller G. The magical number 7 plus or minus 2: some limits on our capacity for process. *Psychological Review*. 1956;63:81–97.
- Miller C, Parasuraman R. Designing for flexible interaction between humans and automation: delegation interfaces for supervisory control. *Human Factors*. 2007;49:57-75.
- Morrow P, Fiore S. Communications in socio-technical systems. In: Lee J, Kirlik A, editors. *The Oxford handbook of cognitive engineering*. Oxford (UK): Oxford University Press; 2013.

- Mott D, Shemanski R, Giammanco C, Braines D. Collaborative human-machine analysis using a controlled natural language: SPIE next-generation analyst III [accessed 2015 Apr]. <https://www.usukitacs.com/node/2852>.
- Murphy RR, Burke JK. The safe human-robot ratio. In: Barnes M, Jentsch F, editors. *Human-robot interactions in future military operations*. Farnham (UK): Ashgate; 2010. p. 31–51.
- Nilsson NJ. *Introduction to machine learning: an early draft of a textbook*. Palo Alto (CA): Stanford University; 1998. p.1.
- O'Dell J. The driverless car: future is almost here; 2013 [accessed 2016 June 17]. www.edmunds.com/cartips&advice/car-technology.
- Ophir-Arbelle R, Oron-Gilad T, Borowsky A, Parmet Y. Is more information better? How dismounted Soldiers use video feed from unmanned vehicles: allocation and information extraction considerations. *Journal of Cognitive Engineering and Decision-making*. 2013;7(1):26–48.
- Oron-Gilad T. Scalable interfaces for operator control units: common display to conduct MOUT operations with multiple video feeds (final research report). Be'er Sheva (Israel): Ben Gurion University of the Negev; 2014.
- Oron-Gilad T, Redden ES, Minkov Y. Robotic displays for dismounted warfighters: a field study. *Journal of Cognitive Engineering and Decision-making (Special Issue on Improving Human-Robot Interaction in Complex Operational Environments: Translating Theory into Practice, Part I)*. 2011;5:29–54.
- Osborn K. Army building smarter robots. Ft. Belvoir (VA): Army ALandT Magazine; 2011 [accessed 2016 June 17]. http://www.usaasc.info/alt_online/article.cfm?iID=1101andaaid=07.
- Parasuraman R, Riley V. Humans and automation: use, misuse, disuse, abuse. *Human Factors*. 1997;39:230–253.
- Parasuraman R, Sheridan TB, Wickens CD. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*. 2000;30(3):286–297.
- Parasuraman R, Miller C. Trust and etiquette in high-criticality automated systems. *Commun ACM*. 2004;47:51–55.
- Parasuraman R., Manzey DH. Complacency and bias in human use of automation: an attentional integration. *Human Factors*. 2010;52(3):381–410.

- Parasuraman R, Barnes M, Cosenzo K. Adaptive automation for human robot teaming in future command and control systems. *International Journal of Command and Control*. 2007;1(2):43–68.
- Pellerin C. Learning systems, autonomy and human-machine teaming. Washington (DC): DOD News; 2015 Nov 13 [accessed 2017 June 16]. <http://www.defense.gov/News-Article-View/628124/work-human-machine-teaming>.
- Pettit RA, Redden ES, Carstens CB, Hooper, D. Scalability of robotic controllers: effects of progressive autonomy on intelligence, surveillance, and reconnaissance robotic tasks. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2013 Jan. Report No.: ARL-TR-6182.
- Pomranky-Hartnett G, Elliott L, Mortimer B, Mort G, Pettitt G, Zets G. Soldier-based assessment of a dual-row tactor display during simultaneous navigational and robot-monitoring tasks. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015 Aug. Report No.: ARL-TR-7397.
- Purdy EM. The increasing role of robots in national security. *Defense AT&L*. 2008;5:26–29.
- Rao A, Georgeff M. BDI agents: from theory to practice. *Proceedings of the First International Conference on Multiagent Systems*; 1995 June; Menlo Park, CA. p. 312–319.
- Russell SJ, Norvig P. *Artificial intelligence: a modern approach*. 3rd ed. Saddle River (NJ): Upper Prentice Hall; 2009.
- Schaefer KE, Evans III, AW, Hill SG. Command and control in network-centric operations: trust and robot autonomy. *Proceedings of the International Command and Control Research Technology Symposium (ICCRTS)*; 2015 June; Annapolis, MD.
- Schaefer KE, Scribner D. Individual differences, trust, and vehicle autonomy: a pilot study. *Proceedings of the 2015 Human Factors and Ergonomics Society*; 2015 Oct 26–30; Los Angeles, CA.
- Scheutz M. The effects of moral competency in autonomous agent architectures. In: Muller V, editor. *Fundamental issues of artificial intelligence*. Berlin (Germany): Springer; 2016 June. p. 517-527.
- Schulte A, Meitinger C. Introducing cognitive and co-operative automation into uninhabited aerial vehicle work systems. In: Barnes M, Jentsch F, editors.

- Human-robot interactions in future military operations. Farnham (UK): Ashgate; 2010. p. 145–171.
- Selkowitz A, Lakhmani S, Larios C, Chen JYC. Agent transparency and the autonomous squad member. Presented at the 2016 International Annual Meeting of the Human Factors and Ergonomics Society; 2016 Sep 19–23; Washington DC.
- Singer PW. War of the machines. *Scientific American*. 2010;303(1):56–63.
- Stone M. Brief of autonomy initiatives to Autonomy Priority Briefing Committee. Wright Patterson Air Force Base (OH): Air Force Research Laboratory (US); 2012.
- Stowers K, Kasdaglis N, Rupp M, Chen J, Barber D, Barnes M. Insights into human-agent teaming: intelligent agent transparency and uncertainty. Presented at the Applied Human Factors and Ergonomics (AHFE) Conference; 2016 July 27–31; Orlando, FL.
- Suantak L, Momen F, Rozenblit J, Barnes M, Fichtl T. Intelligent decision support of support and stability (SASO) operations through symbolic visualization. *Proceedings of the IEEE International Conference of Man, Systems and Cybernetics*. 2001;5:2927–2931.
- Tiron R. Preview: lack of autonomy hampering progress of battlefield robots. *National Defense*. 2003;87(594):35.
- Van Erp J. Tactile displays for navigation and orientation: perception and behavior. Leiden (The Netherlands): Mostert and Van Onderen; 2007.
- Wang N, Pynadath D, Hill S. Trust calibration within a human-robot team: comparing automatically generated explanations. *Proceeding of the 11th ACM/IEEE International Conference on Human Robot Interaction*; 2016 Mar 7–10; New Zealand. p. 109–116.
- Whiten A, editor. *Natural theories of mind: evolution, development and simulation of everyday mindreading*. Oxford (UK): Basil Blackwell; 1991.
- Wickens CD. Designing for situation awareness and trust in automation. *Proceedings of the International Federation of Automatic Control (IFAC) Conference*; 1994 Sep 27–29; Baden–Baden, Germany. Amsterdam (The Netherlands): Elsevier Ltd.; 1994. p. 174–179.
- Wright J, Chen JYC, Quinn S, Barnes M. The effects of level of autonomy on human-agent teaming for multiple-robot control and local security

maintenance. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2014 Apr. Report No.: ARL-TR-6905.

Wright JL, Chen JYC, Barnes MJ, Hancock PA. The effect of agent reasoning transparency on automation bias: an analysis of performance and decision time. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2016 in press.

Xue P, Mott D, Giammanco C, Copestake A. An approach to handling linguistic and conceptual difference across coalition partners. Presented at the International Technology Alliance in Network and Information Sciences Annual Fall Meeting; 2015 Sep; College Park, MD.

List of Symbols, Abbreviations, and Acronyms

ARL	US Army Research Laboratory
ARPI	Autonomy Research Pilot Initiative
ASM	Autonomous Squad Member
CE	Controlled English
COA	course of action
DOD	US Department of Defense
DSB	Defense Science Board
IA	intelligent (nonhuman) agent
IDF	Israeli Defense Force
IMPACT	Intelligent Multi-UxV Planner with Adaptive Collaborative/Control Technologies
L	level
ML	machine learning
NLP	natural language processing
SA	situation awareness
SAT	SA-based Agent Transparency
U	uncertainty
UAV	unmanned aerial vehicle
UGV	unmanned ground vehicle
UV	unmanned vehicle

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIRECTOR
(PDF) US ARMY RESEARCH LAB
RDRL CIO L
IMAL HRA MAIL & RECORDS
MGMT

1 GOVT PRINTG OFC
(PDF) A MALHOTRA

1 DIR USARL
(PDF) RDRL HRB DE
M J BARNES

INTENTIONALLY LEFT BLANK.