ARL

US Army Research Laboratory

Defining "Soldier Intent" in a Human–Robot Natural Language Interaction Context

by Eric Holder

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.

ARL

US Army Research Laboratory

Defining "Soldier Intent" in a Human–Robot Natural Language Interaction Context

by Eric Holder
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)*	2. REPORT TYPE	3. DATES COVERED (From - To)
October 2017	Technical Report	September 15, 2016–August 15, 2017

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Defining "Soldier Intent" in a Human–Robot Natural Language Interaction Context	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Eric Holder	R.0022512.30.1
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
US Army Research Laboratory Human Research and Engineering Directorate (ATTN: RDRL-HRB-DE) 2800 Powder Mill Road Adelphi, MD 20783-1138	ARL-TR-8195

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
The work described was conducted to support a US Army Research Laboratory Director's Strategic Initiative research program to examine intent as conveyed and understood between humans and machines and used to achieve cooperative goals (i.e., "teaming"). The goal of this report is to analyze the question of what intent means for this paradigm, considering the Army context(s) of intended use as well as users' expectations and a range of potential meanings and applications of intent-based behaviors. The results are a review of the relevant literature combined with input from the field and integrated into a consolidated framework. This framework identifies a range of types of intent behaviors scaled for complexity primarily on the need for increased interpretation and reasoning and the use of implicit/existing knowledge (not directly or immediately provided) to ensure understanding and compliance. This framework will provide a structure, a goal, and benchmarks for existing and future efforts.

15. SUBJECT TERMS
intent, human–robot interaction, human factors, autonomy, Natural Language Dialog, shared understanding

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	62	Eric Holder
					19b. TELEPHONE NUMBER (Include area code)
Unclassified	Unclassified	Unclassified			(520) 538-4701

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

INTENTIONALLY LEFT BLANK.

1. Introduction

1.1 Purpose

The work described was conducted to support a US Army Research Laboratory (ARL) Director's Strategic Initiative (DSI) research program to examine intent as conveyed and understood between humans and machines and used to achieve cooperative goals (i.e., "teaming"). The goal of this report is to analyze the question of what intent means for this paradigm, considering the Army context(s) of intended use as well as users' expectations and a range of potential meanings and applications of intent-based behaviors. The results are a review of the relevant literature combined with input from the field and integrated into a consolidated framework. This framework identifies a range of types of intent behaviors scaled for complexity primarily on the need for increased interpretation and reasoning and the use of implicit/existing knowledge (not directly or immediately provided) to ensure understanding and compliance. This framework will provide a structure, a goal, and benchmarks for existing and future efforts.

Note that this project was focused on using natural language exchanges between the human and robot[*] and this impacted, or framed, the analysis and scope. However, communication is impossible if the human and an intelligent agent such as a robot have world models that are out of alignment. Certainly a shared intent-specific task is necessary, but generally in military environments an understanding of intent between operators and their chain of command, as well as the interaction between operators and agents, needs to be understood in their context of use. In particular, it is essential to understand doctrinal concepts such as Commander's Intent (CI); how basic intent information is supplemented via orders and other methods; and how network-centric changes to operations, synchronization, and directive control ensure efficient human–agent interaction in military environments (Barnes et al. 2017).

1.2 Background

As autonomous systems become more commonplace, it is increasingly important that they have the ability to be able to unobtrusively reason about and adapt to their human teammates. However, compared to the ease and flexibility with which humans cooperate with each other, current methods humans use to direct autonomous systems are either slow and labor intensive or extremely limited in

---

[*] The terms robot and machine will be used interchangeably for the rest of this report. Robots are also considered to include unmanned aerial vehicles and unmanned ground vehicles.

scope. This situation effectively reduces human–robot teams to teams of humans operating robots. To alleviate this burden on the human and expand the scope of human–robot operations, this project investigates fundamental research issues in the autonomous inference of human intent and the use of inferred intent to achieve cooperative goals. There is a consistent push for changing the ratio of multiple humans controlling one robot, or unmanned asset, to one human in charge of multiple robots or unmanned assets for military and other operations (Bray-Miners et al. 2012), even to the point of human and swarm* interaction (Arquilla and Ronfeldt 2000; Crandall et al. 2017). As the military environment increases in complexity, it will not be possible for a human to carry out detailed control and monitoring of multiple assets (Chen and Barnes 2014). Therefore, concepts such as intent-based command and control relationships between human and robot teammates may prove essential for supporting these future paradigms.

2.   Intent in Human-to-Human Interactions

2.1 Information Content and Understanding

The first area of focus was looking at what could be learned from human-to-human interactions concerning intent, asking the question of what could then be applied to human–robot collaborations. The human-to-human aspect will have a large role in defining the expectations of the human user in these interactions and can help to structure the input, interpretation, and outputs required to ensure effective human–robot interactions (HRIs). The dictionary definition of intent is "the aim/purpose . . . the thing that you plan to do or achieve" (Merriam-Webster). As opposed to a plan, it implies flexibility on how to accomplish a desired end state.

2.1.1 Commander's Intent

However, the phase "Soldier intent" implies a more specific process in a military context. The term will be interpreted in the light of the well-ingrained concept of CI. As Shattuck (2000) notes, "The concept of intent is written into our doctrine and taught in our schools". The CI is a written statement included in the orders for a mission and should succinctly describe what constitutes success for the operation (see Fig. 1 for an example CI). The orders will include the intent from 2 echelons above as a way to synchronize efforts and are used by the recipient to develop their

---

* Swarms refer to collective robot behavior, often modeled after biological swarms (e.g., ants and bees), that typically have no centralized coordination mechanisms behind the synchronized operation of individual actors (insects), yet their system-level functioning is robust, flexible, and scalable.

courses of action (COA). As noted in Army Doctrine Reference Publication (ADRP) 5-0, Section 2-92, the statement should include the following:

> The operation's purpose, key tasks, and the conditions that define the end state. It links the mission, concept of operations, and tasks to subordinate units. A clear CI facilitates a shared understanding and focuses on the overall conditions that represent mission accomplishment (DOA 2012).

The purpose of the CI is to provide subordinates with guidance yet still allow them the flexibility to accomplish the mission, given the often dynamic nature of the modern battlefield. This is referred to as disciplined initiative (Dempsey and Chavous 2013). Referencing the CI the Soldier can plan and also adapt to the situation, asking him or herself if the new plan still supports the overall intent of the commander. In this respect it is analogous to matching to a higher-order goal or guidance, adapting if a specific path is not available, to a path that can still accomplish the higher goal. For human use, it is recommended to make this intent concise so that it is easy to remember and can be clearly understood by subordinates 2 levels below. The objective of writing a good intent statement is to make it detailed enough to be useful yet general enough so as not to unnecessarily constrain actions.
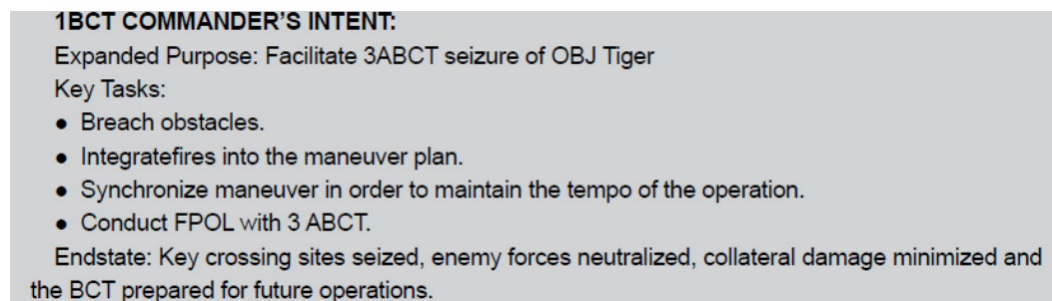
1BCT COMMANDER'S INTENT:
Expanded Purpose: Facilitate 3ABCT seizure of OBJ Tiger
Key Tasks:
- Breach obstacles.
- Integratefires into the maneuver plan.
- Synchronize maneuver in order to maintain the tempo of the operation.
- Conduct FPOL with 3 ABCT.
Endstate: Key crossing sites seized, enemy forces neutralized, collateral damage minimized and the BCT prepared for future operations.

Fig. 1   Example CI taken from Dempsey and Chavous (2013)

Reviews of the effectiveness of the CI process have provided some insights and overall shown that even between humans there is still room for misunderstanding. In particular, an intent statement alone without additional background knowledge will not be sufficient. These results can be considered in the 4 critical components of sharing intent that Shattuck (2000) identifies: formulation, communication, interpretation, and implementation. Shattuck argues that all pieces are equally important but the bulk of the effort in training and doctrine is focused on formulation. These 4 components will also be key considerations when discussing and designing for the HRI context.

Klein (1993, as reported in Shattuck 1995) collected 97 intent statements for analysis and found a wide variability in the content and structure. Their lengths

ranged from 21 to 484 words, with most of them averaging between 76 and 200 words. Further, the analysis found that intent statements often did not comply with the doctrine's content and structural guidance; again showing the difficulty of conveying intent even in well-structured environments.

Murphy examined intent in an analysis for the Australian Defence Science and Technology Office (2002). He found differences in the components of intent that were prioritized between US and Australian Soldiers. Differences by echelon were also found, for instance in the importance of interpersonal contact for the dissemination of intent, specifically if Soldiers felt intent had to be conveyed face-to-face as opposed to in written orders. He also found the complexity of the scenario impacted understanding and that results varied across scenarios, and even for teams within a scenario. His conclusion was that shared intent is an outcome of many variables including situational issues, professional knowledge, and the complexity of the scenario.

In a more formal study, Winner et al. (2007) attempted to devise a quantitative metric for the Shared Interpretation of Commander's Intent (SICI). The authors' background research indicated that although doctrine states that the "tasks" in the intent statements are not to be tied to specific COAs (e.g., prescribe COAs), this is often violated in practice. Their experimental manipulation was to compare conditions where the CI listed specific COAs as opposed to conditions listing priorities in terms of risk, cost, and so on as criteria for COA selection (e.g., prescribed versus inferred COAs). The results supported the hypothesis that making the values by which actions are to be prioritized more explicit enhances both shared interpretation and adaptability compared to statements that prescribe command preference for specific COAs. Transparent priorities and rules may present an effective way to guide COA choices for the robotic teammates while providing situational awareness to the human operator. Discussions with Soldiers suggest they can easily relate to priorities such as maintaining contact with the enemy versus rescuing a "battle buddy" when the latter could endanger additional Soldiers and the mission (Intelligence Sergeant, Fort Huachuca, Arizona, 2017). The use of rules, such as rules of engagement (ROE), are also well understood in the military. Work by Spiker et al. (2007) argues that these types of "if X then Y" rules are what allow Army teams to engage in taskwork without having explicit communications, as the actor knows what to do when encountering event X and the teammate also knows what the actors will do in this situation.

In a field experiment, Shattuck (2000) evaluated CIs by having battalion commanders create operational orders that included CI statements and then the subordinate company commanders used those to create their own orders. During the exercise, 2 situation reports (SITREPs) were created by an external third party

based on the battalion and company orders. The first SITREP informed the participant of changes to the situation that blocked the company from completing their specific mission but could still allow completion of the higher-order objectives (intent). The second SITREP informed the participant that the companies had completed their mission easily and early and then had to decide what to do next (implied to support the CI). Both exercised the company's ability to respond within intent. The battalion commanders were asked to identify what they thought the companies would do in each scenario based on the orders and the intent they had personally written. The company commanders then got the same SITREPS and were asked how they would respond. The battalion commanders were then asked to judge the company commanders' responses. Out of the 32 possible episodes only 17 (52%) of the company commanders' responses were considered by the battalion commander to be a match to their intent. More detailed examination of the 17 matches showed that 6 were not very good matches, suggesting only a 32% match rate. What is really interesting is that the time the company and battalion commanders had worked together did not have a significant impact on the results. Another interesting set of results were the cases where the battalion and company commanders disagreed on the meaning of key terms, such as "delay", which impacted effective decision making. Both of these results will be revisited later when discussing implicit components of intent. Shattuck refers to the communication of intent as the means by which remote supervisors (commanders) impart their presence to local actors (subordinates) given the well-structured but dispersed chain of command in the military. Shattuck takes the argument one step further to note that to incorporate social norms, expectations, trust, or intimate personal knowledge of subordinates, the commander has to also impart his or her presence. This construct is similar to what will be addressed later in the discussion of implicit components of common intent.

2.1.1.1 Network-centric Applications

Neely (2003) explores how CI will have to adapt for network-centric warfare in his concept of network-centric CI. Network-centric warfare is seen as a revolution in military affairs and is defined as follows:

> leveraging computer and communication networks to pervasively link dispersed military forces, providing nearly instantaneous command and control (C2), along with a shared awareness of the battlespace, allowing our military forces to more effectively and efficiently fight and win wars (p. 5).

The variables that are seen to change as CI modernizes for network-centric warfare are (p. 16):

- connectivity between the commander and subordinates

- dispersion and decentralization of assigned forces

- level of shared situational awareness (common operating procedure—COP)

- ability for bottom-up self-synchronization of forces

- speed of command

The primary changes Neely sees are a more static, long-term CI to cover a whole operation and gaps in connectivity augmented by more dynamically updated shorter-term CI posted to the COP and capitalizing on the capabilities of network-centric warfare.

2.1.1.2 CI Summary

The core take-away from the CI analysis to the human–robot collaboration context is that for robots to show true intent-based behaviors the key is for them to be able to adapt, or flexibly perform, in a dynamic environment. This adaptability, or "disciplined initiative", is described by Bryant et al. (2001) as follows:

> "disciplined initiative" is the core benefit, and likely expectation, of truly intent-based robotic behavior for the intended user group. These characteristics will differentiate between understanding an intent and just following commands provided. These aspects can be considered within the terms "Freedom of Action" and "Expected Initiative".

These results also make it clear that an intent statement alone is not enough for complete understanding and mission completion for humans and will also not be sufficient for robots to complete most tasks. The question is how much additional information is required for this understanding and how this information should be provided, as well as when (with the intent or as preprovided underlying knowledge).

2.1.2 Information Required Beyond the CI

2.1.2.1 Information from Orders

Much of this additional information in the military context will be contained in the other sections of the standard 5-paragraph mission order: Situation, Mission, Execution, Administration and Logistics, and Command and Signal (USMC 2015), each providing specific information to help fill in the interpretation and understanding of the operation and the intent.

Dempsey and Chavous (2013) argue that the Concept of Operations, a component of the execution paragraph described previously, is one of the most critical pieces as it is the only element of an order where a commander communicates how all of their forces will combine efforts to accomplish the mission and it tells the story of the battle. They argue the following:

> it should cover the type of offensive, defensive, reconnaissance, or security operation; describe forms of maneuver; identify formations; describe actions on contact; describe the timing of the operation; define the cooperation between maneuver forces at critical points in the fight; and describe how all arms will be coordinated (p. 63).

It should be considered how this or similar concepts could, or should, contain or provide information to guide the human and robot coordination as well.

There is a large volume of research on creating effective orders but the goal here is not to review that but rather to consider what is the essential information that a robot or machine would need from these orders and other sources to correctly interpret and implement the intent of a human tasking them. There is not likely to be a set list of information items for every situation and context and what the robot will need to know will largely depend on its specific tasking, role, and capabilities. For instance, a robot with object detection capabilities will require the information that allows it to identify key objects (e.g., enemy, friendly, weapons, other threats, and entryways) relevant to a mission and would benefit from knowledge of likely locations. For obstacle avoidance and concealment the robots will have to understand the terrain and its affordances. If the robot is expected to coordinate actions with the humans then the timing of specific events and the indicators of events initiating or changing phases will be essential. It is quite feasible in the future that a robot can be provided the documents in electronic form and extract this information directly. It could also be possible to provide all of the information that is electronically available and the robot utilizes only what is relevant. This paradigm would be dependent upon the robot being able to successfully make that determination autonomously or with human intervention or support. The biggest

risk is that the preparation and planning for the robot teammate requires so much additional workload from the human that it is counterproductive. A partial solution is to constrain the robot's lexicon and inference engines to situations the robot is likely to encounter in its mission environment (Barnes et al. 2017).

2.1.2.2   Commander's Vision

Builder et al. (1999) researched the idea of a command concept, "a commander's vision of a military operation that informs the making of command decisions during that operation" (p. iii). They asked the question of "what would the commander have had to tell his/her subordinates before the battle in order to have made their subsequent actions conform to his concept?" (p. 21). They argue that the ideal is only to provide the minimal information that helps a commander convey his command concept, or alter it. Builder et al. (1999, p. xv) list the elements that should be found in an ideal command concept:

1) Time scales that reveal adequate preparation and readiness, not just of the concept but of the armed forces tasked with carrying out that concept.

2) Awareness of the key physical, geographical, and meteorological features of the battle space—situational awareness—that will enable the concept to be realized.

3) A structuring of forces consistent with the battle tasks to be accomplished.

4) Congruence of the concept with the means for conducting the battle.

5) What is to be accomplished, from the highest to the lowest levels of command.

6) Intelligence on what the enemy is expected to do, including the confirming and refuting signs to be looked for throughout the coming engagement.

7) What the enemy is trying to accomplish, not just his capabilities and dispositions.

8) What the concept-originating commander and his forces should be able to do and how to do it, with all of the problems and opportunities—not just the required deployments, logistics, and schedules, but the nature of the clashes and what to expect in the confusion of battle.

9) Indicators of the failure of, or flaws in, the command concept and ways of identifying and communicating information that could change or cancel the concept.

10) A contingency plan in the event of failure of the concept and the resulting operation.

Builder et al. (1999) also note that a complete command concept is rare, if not nonexistent, in practice. Further thought needs to be given to ensure that the information provided is useful for a specific context and the intended user, either human or intelligent agent. A useful parallel would be to consider the items identified by Spiker et al. (2007) to provide shared understanding for Army teams. The key information was broken down by situation, mission, and team as shown in Fig. 2. This approach was seen as a way to make more explicit, and tangible for measurement and Soldier understanding, abstract concepts such as team mental models. For instance, those from the shared mental model work focusing on task, team, and equipment mental models as proposed by Cannon-Bowers et al. (1990). The list by Spiker et al. (2007) provides a good starting point for considering what needs to be known or shared to keep teammates on the "same page". What needs to be examined further is if the information will be usable to a robot in the same form and content as used by humans, or if the information will need to be augmented or interpreted further using such instruments as predicate rules or including robot-specific implications or coding. In general, it needs to be determined how much and what type of knowledge needs to be instantiated in intelligent systems to ensure mission completion and flexibility. The workflow of such processes needs to be considered carefully. As discussed in Section 3, there could be a range of intent-based behavior that will need various information depending on factors, such as allowed initiative and the team and command structure between humans and robots (e.g., 1 on 1 or multiple interacting partners) among others.
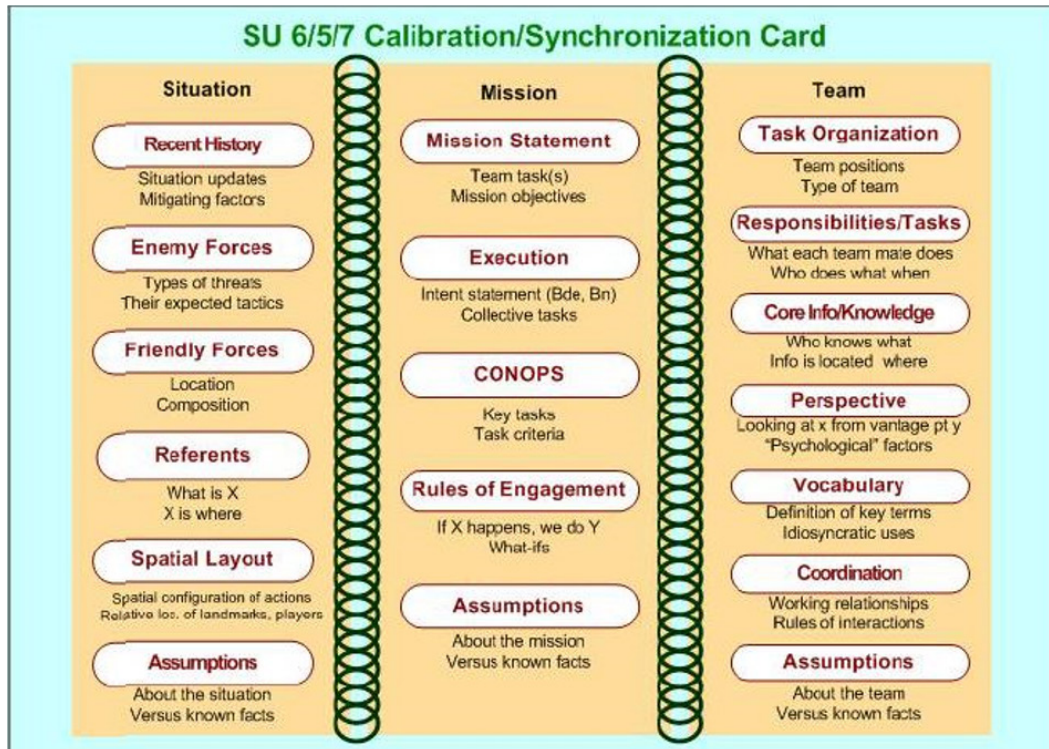
Fig. 2   Shared understanding (SU) calibration/synchronization card (Spiker et al. 2007)

2.1.2.3   Contextual Differences

The context of use, for instance planning versus action phases, can also have a big impact on the range of behaviors and interactions required. The planning stage is seen as the time where all the base knowledge will be provided and the plan or goals exchanged and confirmed. At this stage there is more time for information exchange and understanding checks. The action, or tactical, phases of the operation will likely require fast reactions to inputs based on the dynamic, unfolding situation. The nature of the information exchange between human and robot will also change in the action phase focusing more on tactical intent behaviors (e.g., look in that window), checkpoints on plan progress (e.g., reached Waypoint 1), indicators of phase (e.g., entered building, stopped moving), or behavioral changes (e.g., shots coming from rooftop) keeping the higher level intent behaviors provided at the planning phase (e.g., overall mission objectives and priorities) in consideration. In more advanced HRIs the tactical phase can also provide cues or feedback that the overall plan, or goal, is no longer valid given the dynamic situation, but this would rely on a solid understanding of these conditions as imparted in the planning phase.

2.1.2.4 Implicit Knowledge (Not Directly Spoken or Written)

The next step is to consider the aspects that are not said, or written, in orders. This is touched upon by those looking at expanded versions of command intent, shared intent, and common intent. This is also tied into the German concept of *Auftragstaktik,* or directive control, as described by Shattuck (2000) along with his own concept of presence. All of these concepts discuss the utilization of additional information for intent inferencing and have some similarities. A general overview and some of the subtle differences will be discussed in the text to follow.

The first differentiation is drawn out by McCann et al. (2002) in their work on *common intent* and distinguishes between explicit and implicit intent and highlights that this understanding is "common" meaning between multiple individuals. Explicit, also termed public intent, when used in the military context would include such things as orders and standard operating procedures that explicitly communicate intentions, such as those previously described. Implicit intent would include expectations, beliefs, and values borne out of personal, military, and cultural experience that are often assumed (Bryant et al. 2001). This implicit intent is the more tacit knowledge that is often used to guide or interpret behaviors. It is further argued that those with less shared implicit intent in common will require more explicit intent to achieve the same level of common intent and performance. This was seen as a potentially greater concern for joint teams that do not share as much background (Bryant et al. 2001).

Bryant et al.'s (2001) review of common intent discusses how the *Auftragstaktik* requires at least some implicit intent through an organization that requires and rewards initiative. It is argued that to do so:

> commanders need more than a directive; they need a lifetime of cultural
> experience to learn how to act on their own initiative, trust that others will
> be doing the same, and an understanding that they will be rewarded rather
> than punished for taking risks. (p. 28)

This is not something that can be created on the spot or in the moment. This also brings up the question of how much could be shared with an intelligent agent in a document or database, such as reading through, versus living within, doctrine or others' written accounts of military experiences.

The discussion by Bryant et al. (2001) continues to examine the shared history between the units that is built up before the start of a mission based on common doctrine and procedural standards, as well as through interactive methods that are in place to build shared understanding and shared intent, such as backbriefs, postorder discussions, and mission rehearsals. One key aspect that is often implicit

is a common language, or terminology, and this has been seen in the problems Shattuck (2000) reported with understanding CI, as well as the work on shared understanding by Spiker et al. (2007). Having a common terminology is seen as key to facilitating the rapid execution of orders without excessive discussion or misinterpretation (Bryant et al. 2001). It is argued that given the natural limitations of any language or vocabulary the interpretation of an explicit communication will always be based on an extensive network of implied meanings that are used to modify the specific words. This is addressed by Clark and colleagues in their constructs of common ground and grounding that were developed as part of their "contribution theory" of conversation, considering not just the utterance itself but how each party seeks to make sure the utterance was understood (Clark et al. 1983; Clark and Brennan 1991). The authors argue that people do not have to explicitly express the majority of information underlying a message. Instead, they rely on the vast knowledge of the addressee to allow comprehension. Common ground is seen to consist of the mutual knowledge, beliefs, and assumptions held by people.

Bryant et al. (2001) describes the implicit intent as deriving from all of the experiences in and out of the military that form beliefs, values, habits, expectations, and personal styles. It is important to note that where explicit intent can be built from explicit communications, implicit intent is primarily built via shared experiences. This begs the question, if this type of implicit intent is deemed essential for HRI, how could such experiences be gained, embedded, or otherwise created? The logical follow-up question is how to differentiate applicable knowledge between one context and the next for robot behavior. Namely, if a specific robot learns from repeated experiences with a specific team in one context should those experiences be generalized to the next context (mission, environment, and task, etc.) or shared to all other robot teammates working with other teams? Realistically, a cognizant robot even for constrained environments is not likely to be available either from the laboratory or off the assembly line in the near future. Fortunately, the problem of not being able to preprogram every eventuality is now widely recognized by the machine learning community. Techniques such as reinforcement learning, neural nets, and genetic algorithms are evolving to enable robots to adapt to complex environments. Understanding will be accrued through iterative programming or using machine learning techniques during multiple field exercises (Sutton and Barto 1998; Wang et al. 2016; Barnes et al. 2017; Chen et al. forthcoming).

Gustavsson et al. (2008) in their work on defining a machine-interpretable format of CI define the implicit content of intent as what is developed over a longer period of time to include "the expressives and the concepts, policies, laws and doctrine agreed to by military, civil, organizations, agencies, nations and coalitions" (pg 4).

They further define expressives as a "component of CI that describes the style of the commander conducting the operations with respect to experience, willingness for risk, use of power and force, diplomacy, ethics, norms, creativity and unorthodox behavior" (pg. 5). The example given for working with a Nordic Battle Group would be making the implicit style of using low violence, implicit to the other nations that have different doctrine, cultures, and backgrounds. This has much overlap with the concept that Shattuck (2000) terms presence. Their overall goal was to make CI less ambiguous and more widely distributed to guide subordinate action and use in simulations and computer systems of multinational forces. Their outputs were structured around effects-based thinking and a proposed Battle Management Language usable by human and computer agents.

It is not foreseeable in the near term for a robot to need, or be able to utilize, much of the implicit knowledge. Initially the robot will make simple adjustments, such as "Obstacles ahead—need to replan". What is required would need to be coded in rules that can be followed (e.g., minimize harm to humans, always give priority to casualty assistance) and ways to map those rules on to the missions and contexts where they apply. This would match on to the idea of augmenting reduced implicit intent with increased explicit intent, being careful to minimize the workload of the human operators responsible for conveying this explicit intent to the robot.

In one final distinction between the various terms used, Bryant et al. (2001) differentiate between *common intent* and *command intent*. The authors define common intent as "the complex organization of knowledge, values, practices, and attitudes of individuals" (p. 119), and then command intent as "the more specific organization of knowledge, principles, attitudes, values, goals, and constraints that make up a commander's plan for a specific operation" (pg. 119). Command intent is seen to be based, in part, on common intent and to serve as a link between common intent and the implementation of C2. Furthermore, command intent is tied to a specific time, place, and operational goal, whereas common intent is related to a lifetime of experiences inside and outside of an organization and across various mission contexts. The focus on implicit intent, if added to the HRI context, should be more on the domain of command intent.

2.1.3 Research Considered Outside of Application Scope

It should be noted that this review did not include work on understanding enemy intent, as this was not seen to be applicable to this context. The context of focus is a collaborative context where the humans and robots are on the same team and can query each other and expect truthful answers. This is a whole different paradigm of intent inferencing than employed in the enemy intent context. Likewise, paradigms such as theory of mind (ToM) were reviewed (Levin et al. 2008) but not included,

as the level and format of mental states, representations, and beliefs, desires, and goals used for the ToM paradigm to analyze or explain behaviors were not seen to be necessary, or applicable, for robots to understand intent and perform the required actions. The relevant overlapping concepts, such as seeing things from another's perspective, are thought to be best coded for robots in their reasoning rules and are covered in the other paradigms previously described. The robots are being tasked in the expected contexts, rather than freely interacting with the human to determine more philosophical concepts. This paradigm would map better on to humans trying to infer intent from the robot partner, but this awareness is felt to be better addressed in the transparency of the HRI as discussed in Barnes et al. (2017)

2.1.4 Summary of Information Content and Understanding

Summary points from this section for potential applicability to intent-based behaviors for HRI in the military context include the following:

- The ultimate goal of intent-based behaviors should be guidance, or tasking, that allows disciplined initiative to allow robots to adapt, or flexibly perform, in the dynamic environments expected by military Soldiers and to fulfill expectations Soldiers will have from experiences with the CI.

- There are often misunderstandings in interpretation of intent and therefore a means to confirm, clarify, or correct the intended behaviors needs to be implemented. The robot needs to know what it is to do and the human needs a way to confirm that the robot initially understands the goal and priorities and can then remain aware of progress/divergence from this as the mission proceeds.

  o See work by Chen et al. (2017) on transparency for human–agent teaming effectiveness.

  o Transparent priorities and rules of engagement may present an effective way to guide the action paths and decisions of robotic teammates while still providing situational awareness to the human operator.

- All 4 components of the intent-transfer process need to be considered: formulation, communication, interpretation, and implementation.

- An intent statement alone without additional background information will not be sufficient to allow the correct behaviors for most situations and therefore needs to be supplemented with additional information.

o There is not a set list of items for every situation and context and what the robot will need to know will largely depend on their specific tasking, role, and capabilities.

o It could be possible to provide all of the information that is electronically available (e.g., via orders, doctrine) and the robot ingests this and utilizes only what is relevant, but only if the robot is able to successfully make that determination. Otherwise, this will require human support.

o Ensuring that the robot has all of the additional information required cannot demand excessive workload from the human operator or it will risk eliminating any value added.

o Implicit information concerning expectations, beliefs, and values is often used in human–human intent inferencing but it is not foreseeable in the near term for a robot to need, or be able to utilize, much of the implicit knowledge. What is required would need to be coded in rules that can be followed.

o It has been argued that for teams a deficit in implicit intent will require more extensive explicit intent.

- Having a common terminology is seen as key to facilitating the rapid execution of orders without excessive discussion or misinterpretation.

Some key items of explicit, mission-based information to provide the robot can include:

- The task and purpose (goal/endstate), including more of the why of an assignment to support adaptation

- A common spatial reference grid shared between humans and robots

- Higher-level mission constraints (e.g., ROE, overall strategic objectives or effects desired)

- Relevant pieces of mission plans and orders (enemy and friendly, terrain, timeline and coordinating instructions, constraints, etc.)

- General military tasks that are always important (causalities, immediate threats, etc.)

- Implications of detected changes to the world model (reported by robots or humans), which can result in a new behavioral plan

- Changes in mission stage or priorities provided primarily from additional human input or location information (e.g., moving from scout/recon stage to movement to contact stage and the related tasks and areas of interest)

- Real-time locations of friendly, enemy, and neutral persons (sensors, reports, blue-force tracker, etc.)

2.2 Relevant Methods for Building or Maintaining Shared Intent

2.2.1 Methods

The next important area of focus is to look at what methods of building and maintaining shared, or common, intent could be leveraged for the human–robot intent context. This will leverage work from the military context and the psychology of teams, as well as from the human and robot interaction domain. This is particularly important because Soldier–robot teams will not be isolated components when implemented on the future battlefield; they must understand the intent of other Soldiers and Soldier–robot teams to be effective collaborators on the battlefield. This will require anticipating each other's actions, as well as communicating intent (Barnes et al. 2017; Chen et al. forthcoming).

Spiker et al. (2007) developed a training package to help military, small-unit action teams stay on the same page. This training included several methods to help synchronize SU between teammates and—although much is generalizable—was tailored for face-to-face situations. From a measurement standpoint, SU was treated as a "state" that has transitory properties, where its emergence is enabled through trainable cognitive skills and can be maintained through trainable monitoring and communication skills. The techniques were broken down into 3 components (definitions are from Spiker et al. 2007, p. 17):

> The first component, *mutual monitoring*, is the basic process by which team members observe or monitor one another's taskwork behaviors, teamwork interactions, and other communications within the team environment while a mission-task is being performed….The primary element of this process entails looking for signs of possible SU breakdown [see Fig. 3], such as a frequent need for repeat communications, long periods of no communication, or indications of unexpected or unusual behaviors from a team member.

The training interface included a task designator that provided a short description and state (e.g., ongoing, complete, deleted, on-hold, delegated, or shared) of a team member's current tasks. This idea could be considered when looking at promoting transparency in the HRI context (Chen et al. 2017). The human will surely need a

means to monitor the behavior of the robots. The robots to some extent will also need to monitor the behavior of the humans to ensure smooth coordination and timely actions.
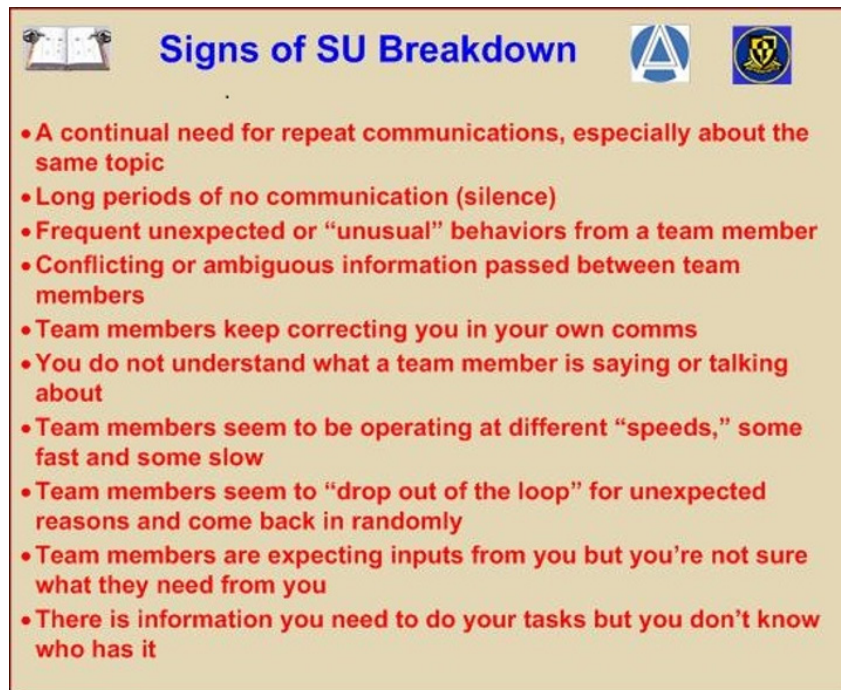

Fig. 3   Signs of breakdowns in shared understanding (Spiker et al. 2007)

The second component, *communication*, is viewed as "a standard process of exchanging information for the purposes of maintaining SU, confirming the status of SU, or restoring a currently degraded state of SU to some acceptable level" (Spiker et al. 2007). For the training scenario tool the format of the communications could include a mix of verbal, email, graphics, extended text messages, face-to-face, or physical observation, among others.

The third component, *intervention*, "entails a more direct means of controlling the synchronization of team mental models through behavioral adjustments (one's own behavior and that of one's team members) and feedback….The intervention might occur in advance of a possible SU breakdown, which is more properly called feedforward, or it could be provided after the fact as classic feedback" (Spiker et al. 2007).

Specific techniques from Spiker et al.'s (2007) 3 component areas that are of particular interest to HRI might include the following:

- Listing and comparing information, tasks, and roles. This is a means of making some of the internal information explicit.

- Tailoring a message for a designated receiver. This entails knowing your audience and what is relevant to them and when.

- Giving a backbrief about something just heard or learned. This is a check on understanding to ensure that the recipient can explain in their own words the gist and main points of the information they just heard.

- Asking critical, focused questions. This is a way of asking questions to ensure core information was understood (e.g., "When will you move your troops?").

- Providing feedback to an individual team member or group of team members. This is an intervention method to correct, or improve on, a specific communication or piece of information a teammate has provided you.

- Engaging in table-top exercises and "what-ifs". This is routinely used in the mission planning process as a way to make possible contingencies and actions explicit.

It should be noted that these techniques can be combined, such as table-top exercises with critical, focused questions and feedback provided during this exercise. One aspect that is mentioned by Spiker et al. (2007) as a performance enhancer is cross-training or being able to see things from a teammate's perspective. Further thought should be given to if, and how, this applies to HRI contexts. For example, Controlled English (CE) is a specialized natural language representation developed by International Technology Alliance scientists from ARL and the United Kingdom. To train the CE inference engines, researchers paired a CE module used for intelligence with an actual intelligence analyst to conduct analytical exercises. The inference rules and algorithms instantiated in the CE artificial agent were based on rules and lexicons learned from interacting with the human analyst (Barnes et al. 2017).

Making the tasks, actions, and priorities of the robot teammate more transparent might provide another surrogate solution. Another aspect to be considered is general knowledge of teams and how they work and if that can, and should, be provided to the robot as explicit information. Further, what type of information would be needed, ranging from teams in general to tailoring to the specific military team it is embedded with, as task expertise will vary across a team, as will the information held and required by each teammate. The capabilities of machine learning will have a big impact on the success of such measures. It also stands to reason that the humans will also need this transparent knowledge of what the robot knows and can contribute to the team to best achieve mission goals. This

information could perhaps take the form of an introduction, or prebriefing, as a unit acquires robot teammates and repeated as the robot adapts or new human teammates are added.

Several of the methods listed by Spiker et al. (2007) are echoed by Shattuck (1995), who lists the methods that commanders used to promote dialogue and impart presence to subordinates in the field. It is noted that time often does not allow all of these methods to be used. The list included the following (pg. 124):

- Briefings: After the orders are written, a face-to-face meeting in which the senior commander explains his or her intent for the upcoming operation and subordinate commanders ask questions to clarify that intent.

- Backbriefs: The senior commander confirms the subordinate's understanding of his or her intent by hearing the subordinate commanders brief their plans.

- Rehearsals: Prior to beginning an operation, the senior commander assembles the subordinate commanders and staff officers to review and rehearse the plan, often maneuvering forces on a terrain representation with an intelligence officer role-playing the enemy.

- Leader's reconnaissance: The senior and subordinate commanders join together to perform a reconnaissance of the terrain on which they will perform the mission and the senior explains his/her vision of how the battle/operation will unfold.

- On-site visits: The senior commander visits subordinate units and confirms their understanding of the upcoming operation.

Rosen et al. (2011) discuss adaptive performance in teams and break down the team performance processes into 3 phases—planning, action, and interpersonal—based on Marks et al.'s (2001) list of 10 dimensions critical to team functioning. Planning activities include transition processes, such as mission analysis, goal specification, and strategy formulation. Action activities, such as coordination, monitoring, and back-up behavior, occur during the periods of time when teams are actively working towards goal accomplishment. Interpersonal activities—conflict management and affect management—can occur in both the transition and action phases. These performance activities are augmented by a situation assessment phase and a team learning phase. The former phases fall more into the earlier discussion of content required and the later phase needs to be considered in terms of human learning and adaptation to the embedded robot partners, as well as what

can be expected or applied in the context of machine learning about team behaviors. Rosen et al. put forward these 4 learning processes for human teams, which have some parallels in machine learning:

> Recap (wherein the team builds understanding of its past performance), reflection (wherein the team diagnoses and evaluates its performance), integration (wherein the team transforms the initial understanding of the team's performance into a new, shared model by incorporating previously identified successes and failures), and action planning (wherein the team purposefully develops lessons learned from the evaluation of its past performance and develops a plan for integrating these lessons into the team) (pgs. 112–113).

The primary question for HRI contexts is if the team learning can/should take place with human and robots together, simultaneously, or for humans and robots independently. These activities related to adaptation are relevant, as a core advantage of intent is to promote, or allow, adaptive behaviors guided by the intent and these behaviors are identified to stimulate teams to coordinate and work better together, promoting emergent states such as mutual trust, motivation, and team situational awareness. It is noted that in dynamic environments with rapid changes teams alter their performance processes on the fly and the robot as part of the team will need to be aware of this.

2.2.2 Communications in Detail

As seen in the methods listed so far, communications play a key role in having a shared understanding of intent. Oron-Gilad (2015) also reports that people seem to favor speech when interacting with robots. Shattuck (1995) reported 6 observed categories of verbalizations in his study in the company commanders' communications when faced with a situation report that reveals a mission anomaly. These can be considered as potential dialogue exchanges to support within Army mission HRI. These verbalizations included the following (p. 86):

- Need for information: Request additional information or state need for additional information to be confident in a decision.

- System status: Statement concerning the status of enemy forces or friendly forces.

- Reference to procedures: Reference to a procedure or other information provided prior to confronting events described in the battalion commander's situation report (e.g., operations order, unit standing operating procedures, doctrine).

- Reference to intent: Reference to the intent of the senior commander or the commander two echelons above.

- Course of action: Description of an action that the company commander would take or believes other commanders would take in response to the situation report.

- Coordination: Description of the efforts the company commander would take to coordinate activities with other units.

There are several other communication-based cues to understanding used in both human-to-human and human-to-machine interactions, but it is important to first consider the medium of the communication and the affordances it allows as these can impact the cues available and the understandings and misunderstandings. Clark and Brennan (1991) describe how different media change the constraints on grounding in conversation. They list 8 features of interaction (p. 141): copresence, visibility, audibility, cotemporality, simultaneity, sequentiality, reviewability, and revisability. Bryant et al. (2001) have adapted this to show how these fit into various example media (see Tables 1 and 2). The focus of the current work is on natural language interchanges but future work will need to consider for each type of information exchange required the media available and the optimal match for required understanding, bandwidth constraints, and mission factors such as distance and heads down time.

Table 1   Classes of constraints on grounding, recreated from Bryant et al. (2001), who adapted from Clark and Brennan (1991)

Constraint	Definition
Copresence	A and B share the same physical environment
Visibility	A and B are visible to each other
Audibility	A and B communicate by speaking
Cotemporality	B receives at roughly the same time A produces
Simultaneity	A and B can send and receive at once and simultaneously
Sequentiality	A's and B's turns cannot get out of sequence
Reviewability	B can review A's messages
Revisability	A can revise messages for B

Table 2   Examples of 7 media and their associated constraints, recreated from Bryant et al. (2001), who adapted from Clark and Brennan (1991)

Medium	Constraint
Face-to-face	Copresence, visibility, audibility, cotemporality, simultaneity, sequentiality
Telephone	Audibility, cotemporality, simultaneity, sequentiality
Video teleconferencing	Visibility, audibility, cotemporality, simultaneity, sequentiality
Terminal teleconferencing	Cotemporality, simultaneity, reviewability
Answering machine	Audibility, reviewability
Electronic mail	Reviewability, revisability
Letters	Reviewability, revisability

The methods introduced by Clark and colleagues to support grounding in communication are of relevance here also as some can be effectively designed into human–robot interactions. Grounding can be seen as the feedback loop from the listener to the speaker to confirm that their utterance was, or was not, heard and understood, establishing the common ground of what is mutually believed. The communication process is seen as 2 phases, presentation and acceptance, with a side sequence for clarification. Grounding is achieved either via positive feedback or negative feedback (e.g., questions, verbal or nonverbal signs of confusion). These are based on the principle of closure, where agents performing an action require evidence, sufficient for current purposes, that they have succeeded in performing it. This can be thought of as feedback on task performance. Some mechanisms of that positive feedback include the following:

- Continued attention: B continues attending to A

- Relevant next contribution: B starts in on next relevant contribution

- Acknowledgment: B nods or says continuer like "uh-huh, yeah," assessment ("That's great!")

- Demonstration: B demonstrates understanding A by paraphrasing or reformulating A's contribution, or by collaboratively completing A's utterance

- Display: B displays verbatim all or part of A's presentation

Koschmann and LeBaron (2003) summarize Clark's methods of accomplishing acceptance as 3 types: a new contribution, an acknowledgement or continuer, or a request for clarification of some or all of the presentation. The authors further argued that many of the principles of common ground and grounding do not always hold up when applied to real-world contexts. They focused on dialogue in an operating room and found that there are often more than 2 people actively involved in a dialogue, silence can play multiple roles, the idea of discrete turns is not often the case, and these "acceptances" may or may not have happened up to 3 "turns" later, often making it ambiguous which "presentation" is being referred to. The environment in real-world interactions is also not static and readily accessible as assumed in Clark's model; in this case the surgeons were manipulating the tissue, which altered references. Further, there is much beyond just the talk, or words, of an interaction (the vocal pairs of presentation and acceptance), such as space, artifacts, and body movements (e.g., pointing) that play an important role as well. The authors also argue that common ground is "a cooperatively constructed mental abstraction" and not really available to anyone. The point of this is that the cues to understanding can be important but the overall context of the interactions need to be incorporated as well to effectively use the cues and further that common ground as a measure in itself may prove to be an intangible concept.

A review by Jurafsky and Martin (2009) dives into examples of how these grounding principles have been successfully and unsuccessfully implemented in various human–machine spoken dialogue systems, mostly for travel arrangements and help systems. One of the key points is that users of these speech-based interfaces are often confused, and performance is degraded, when the system does not provide a clear acknowledgement signal. Also included in the review are the maxims put forward by Grice (1975, as cited in Jurafsky and Martin 2009) that allow hearers to draw correct inferences. These maxims play a guiding role in the interpretation of verbal communication utterances and would be expected to be implemented in any of the systems designed to use natural language to share or clarify intent. The maxims include the following:

1) Relevance: Be relevant. Things are mentioned for a reason.

2) Quantity: Do not make your contribution more or less informative than required.

3) Quality: Try to make your contribution one that is true (do not say things that are false or for which you lack adequate evidence).

4) Manner: Avoid ambiguity and obscurity; be brief and orderly.

2.2.3 Observation of Compliance

Another tangible indicator (cue) that the intent was understood and is being complied with is when behaviors matching that intent are either observed or reported. One such measure was used by Entin et al. (2005) where they computed anticipation ratios, which were the number of information transfers divided by the number of information requests, as a measure of the extent to which each team member anticipated the information needs of other team members. Other measures can include providing backup to a team member when needed (Spiker et al. 2007), reporting at the correct times (e.g., objective achieved), and reports that indicate understanding of the overall mission (e.g., a large group of enemy fighters seen on the route that might block progress). Supporting this mutual monitoring (Spiker et al. 2007) via multiple means, such as direct observation, audible or textual reporting, graphical viewing of progress or camera view, etc. can be tailored to the situation and constraints to provide the optimal HRI.

2.2.4 Summary of Methods

There are several techniques and cues that can be employed in real time or built into the intelligent agent or standard operating procedures to help ensure the shared intent is understood and maintained throughout the mission. These fall into various ways to monitor, communicate, and intervene to maintain this shared understanding, as well as structuring the human–robot interactions in a manner to maximize understanding and minimize misunderstanding, such as leveraging known cues and principles inherent in human-to-human dialogue situations.

2.3 Conclusions on Human-to-Human Intent

Developing an understanding of intent is crucial for creating an agent architecture that interacts with humans in a complex environment. Such dyadic interactions are predicated on communication among humans and agents (e.g., robots) requiring techniques such as sophisticated Natural Language Processing (NLP). Such an interaction depends on mutual understanding of the operational context, which in turn requires understanding of doctrinal underpinnings of the term "Soldier's Intent" including CI, commander's vision, synchronization with other units, and explicit and implicit knowledge components. The discussion indicates that intent in a military context is not a specific COA but rather a deeper understanding enabling the robot (or other agent) to adapt to volatile combat situations. Machine learning, bidirectional transparency, situation understanding, training procedures, tactics, and other measures are discussed as enablers of instantiating and maintaining shared intent within a software architecture that maximizes mutual understanding, relying on NLP, and mixed-initiative decision making.

3. Defining a Range of Intent-Based Behaviors

3.1 Background and Definition

This section describes the range of intent-based behaviors, scaled for complexity primarily on the need for increased interpretation and reasoning and the use of implicit/existing knowledge (not directly or immediately provided) to ensure understanding and compliance. This product was designed for human–robot tasking interactions for military applications in a spoken natural language command context and identifies this and other assumptions of the product in the assumptions subsection. The goal of this taxonomy is to promote discussion of what intent could mean and how to support it, as well as provide a benchmark of sorts to steer development efforts towards it. It is also possible that with real-world applications items higher on the list will prove to require more complexity for a given development context than those items listed later.

Working Definition of Intent: Intent for this work is defined as what the human wants the robot to do and in certain circumstances the ability of the robot (or robotic reasoning component) to infer the human's objective(s) or purpose in order to select behaviors that align.

3.2 Textual Descriptions of Types of Intent Behaviors

A range of 7 types of intent behaviors are described in Table 3 and in the text as follows.

Type 1, Basic Commands: The robot is instructed on what to do, e.g., "Carry this object from point A to point B". There is very little interpretation of the basic command, freedom of action, or expected initiative. Where the robot will have some flexibility is in determining how to proceed based on its own capabilities and constraints. This could for instance be in determining the optimal path between the 2 points to avoid obstacles or hindrances. The information that the robot must have before acting includes a list of commands (playlist) of which one is selected (e.g., relocate from point to point), a planned route, and a spatial reference for the world and locations.

Type 2, Ambiguous References: The robot will have to decipher references to objects and locations to correctly act (e.g., "Go to/behind/near parked car or over there"). These kinds of references can contain ambiguity in both the object (which car) and the location reference (behind, from whose perspective). Additional information from the human such as eye gaze or gestures can be incorporated to clarify references (e.g., where the human is looking or pointing). The interaction

can also leverage knowledge from prior interactions to help clarify pronouns and references as is done in human-to-human interactions. For example, asking "Are you here yet?" implies a previously agreed on meeting place. To be able to perform these behaviors the robot will need to know what objects are, either via object recognition capabilities or by having predefined objects and references in its internal geo-spatial model. The robot will also require the reasoning ability to determine intended spatial references from the correct perspective, given the context.

Type 3, Performance Conditions: The robot will have to interpret more abstract performance conditions and constructs (e.g., "Go there covertly/undetected" and how to implement those in the often dynamic mission context). There can be a range of complexity proceeding from basic rules, such as staying closer to buildings or out of open areas, to more advanced reasoning about the relationships of objects in the world (e.g., affordances such as cannot see through) to these conditional properties (e.g., being seen by whom and lines of sight). To be able to perform these behaviors the robot will need to know how those conditions (e.g., urgent, undetected, self-preservation) relate to possible behaviors. The knowledge can include terrain reasoning as defined by the National Research Council (2002) as the following:

> The ability to use information about natural terrain features (elevation, vegetation, rocks, water), manmade features (roads, buildings, bridges), obstacles (mines, barriers), and weather [to support military maneuver reasoning] using terrain reasoning, mission, friendly and enemy locations to determine the best maneuver and selection of positions for stealth and to support mission package needs (e.g., hull down for direct fire, clear of overhead obstructions for indirect fire) (pg. 19).

To this point the robots are still being told exactly what to do (given their task) with some flexibility in determining how to do it. The following examples move from this explicit tasking to include reasoning about higher-level goals or mission or rules of behavior.

Type 4, Determine Tasking: The robot will have to interpret a higher-level goal or assignment into the specific objects and actions of interest (e.g., "Provide preliminary route clearance before we depart"). Advanced reporting could include providing findings in the perspective that best suits our own force's human performance in that context (e.g., ego or exo-centric, grid reference for targeting, above to your 3 o'clock). It is important for the robot to know what the human wants to know and the different ways that this might be captured by available sensors. To be able to perform these actions the robot will have to have a pre-

existing knowledge base of objects and their properties (e.g., what might be a threat in the given mission context or area of operations). The robot will also need a pre-existing knowledge base of actions that can support various goals/missions/assignments, such as where to look/listen for threats in this context (e.g., windows of buildings, roofs, alleys, etc.). The robot may also require knowledge that allows structuring their task planning, such as priorities (e.g., dwell time, search patterns) for areas and objects in this context. For example, how sure (e.g., 70%, 90%) does a robot need to be that a building or alley is clear before moving on.

Type 5, Adaptation from Goal/Assignment: The robot will be able to adapt its assignment or mission based on changes to the situation. The core to this item is the ability to recognize that the situation has changed in a way that requires adaptation. One instance of this would be the ability to recognize the need to break from a current task, and complete a not-explicitly assigned, higher-priority task (e.g., recover own force casualty, track a high-value target that was identified, respond to an immediate danger to robot or own forces). A second instance would be the robot adapting its tasking to the phases of a mission. For instance, covertly scouting ahead to identify threats before the unit reaches the hostile location, or moves out; and then providing overwatch, responding to locations of sniper fire, or identifying clear avenues to move as forces enter the area; and then providing a lookout as well as monitoring movement inside a building as it is breached, along with monitoring the enemy fleeing the building. To be able to perform these actions the robot will have to have knowledge of a priority hierarchy of tasks for the mission context, as well as always important tasks not listed for each mission. The robot will also need knowledge of the flow, tempo, or stages of a mission, as well as the indicators that a stage or phase has been reached (e.g., command given, troops' location/movement, set time points). The robot will also need knowledge of the planned mission and endstate (such as provided in the 5-paragraph mission orders) and related properties to be able to determine progress towards or prevention in achieving goals. The robot will also need to understand key environmental aspects of the mission and battlespace such as terrain* and decisive actions/events/points. This could also include performing battle damage assessments or determinations when key entities (e.g., enemies, clear path) are not where anticipated, invalidating the overall plan.

Type 6, Multihuman Context: The robot will need to be able to take input and provide information to support mission completion from various humans. This

---

* Terrain is analyzed using the 5 military aspects of terrain—OAKOC: **o**bservation and fields of fire, **a**venues of approach, **k**ey and decisive terrain, **o**bstacles, and **c**over and concealment (DOA 2016)

would be a true team context of HRI where inputs and outputs are flowing between multiple parties rather than one human assigned to interact with the robot. This is especially tricky in a natural language context, as compared to various team members using a graphical user interface (GUI). What makes this complicated is that input and references are no longer constrained to a one-on-one context, which goes beyond the standard grounding context and can increase confusion as was found in the real-world examples used by Koschmann and LeBaron (2003). In this case, acknowledgements and references to prior discussions or different perspectives can be ambiguous for both humans and robots. For instance, is the human that just spoke referring to an object mentioned by another human, an object he or she mentioned earlier, an object he or she is currently looking at, etc.? The robot could also achieve more advanced teamwork behaviors by supporting various team members selectively with the information or actions that the teammate needs to do his or her specific tasks/jobs. To be able to perform these actions the robot will need to be able to recognize what is valid input on mission status or intent from multiple humans (sources); in essence, a filter of what is relevant to guide robot behavior. The robot will also need knowledge of the composition of the mission team and their roles and tasks to determine how to coordinate and communicate effectively. Some of this information could be provided in advance for generic mission teams using doctrine and other publications combined with machine learning. The robot will also need knowledge or rules to determine which inputs to prioritize over others, given the potential for conflicting requests or commands. This type of intent is in some respects a different thread in the progression; as the robots dealing with multiple humans can exhibit behaviors of all the other types listed, it is specifically the multihuman context that sets this type apart.

Type 7, No Consistent Human Input: The robot will be provided a mission-level task or role based on a desired endstate or objectives and can be left on autonomous duty to reason which tasks from a set of possible tasks it could best perform to support/accomplish that mission given the current circumstances. Some of these types of assignments could be conceived as effects-based operations, using the end game as a starting point. This does not rule out occasional human input, but it is not required for the robot to keep performing—meaning it has a very high neglect tolerance as defined by Goodrich and Schultz (2007) as the time between required inputs. This also does not rule out the need for the robot to produce information that is usable by humans on a more constant basis. To increase the utility for extended mission deployment, the robot could integrate additional reasoning to collect activity-based intelligence to determine trends and indicators, as might be needed for Intelligence Surveillance and Reconnaissance activities (e.g., identify patterns of life or atmospherics—indicators of stability, pro/anti-US sentiment, baseline movement and activity patterns, meeting places, traffic, timing of activities,

coherent change detection) (DOD 2011). To be able to perform these actions the robot will need flexible mission tasking similar to CI (goals, endstate, and mission tasks). The robot will also need a model to determine risk and rewards for various actions. The robot will also likely need the ability to assess the situation to determine self-preservation and protection behaviors. It should be noted that autonomous agents are not expected to have authorization to engage in lethal force. If this situation was to change then the reasoning requirements will increase significantly to include determining the priority of specific targets, potential for collateral damage, priority of capture versus kill, the impact on hearts and minds of the civilian population, and other very complicated tradeoffs.

3.3 Table of Types of Intent Behaviors

The types of intent behaviors previously described are presented in tabular form in Table 3 to allow a more succinct comparison. As the reader proceeds through the table the types of intent behavior move from direct taskings to reasoning about higher-level goals, missions, or rules of behavior. The table includes a description of the type of intent, variability to describe variations or additions to the basic concept, and preconditions that list some of the required building blocks to achieve these types of behaviors.

Table 3   Range of intent complexity

Description	Variability	Preconditions
1. Basic commands. Do as I say: "Go to Waypoint 1 and watch"	Plus: Robot can determine some of the how to proceed based on its own capabilities and constraints (e.g., the exact path)	Playlist, planned route, spatial reference for world and location.
2. Ambiguous references. Must decipher objects and references to correctly act "Go to/behind/near parked car"	Plus: Can incorporate other information from human such as gesture or eye gaze to help determine references Plus: Could reference knowledge base to learn from past interaction/dialogue to clarify pronouns and references	• Object identification ability: either predefined objects in world model or object recognition capability • Ability to determine intended spatial references in world model from correct perspective (behind from which perspective) for the given context
3. Performance conditions. Interpret more abstract conditions and constructs for performance such as "Go there covertly/ undetected"	Plus: Proceed from basic rules like "Stay close to buildings and out of open areas" to reasoning about relationships of objects in the world to these properties (e.g., being seen and by whom, lines of sight)	• Terrain reasoning—The ability to use information about natural terrain features (elevation, vegetation, rocks, water), manmade features (roads, buildings, bridges), and obstacles (mines, barriers), and weather.[a] • Military maneuver—Using terrain reasoning, mission, friendly and enemy locations to determine the best maneuver and selection of positions for stealth and to support mission package needs (e.g., hull down for direct fire, clear of overhead obstructions for indirect fire)[a].
4. Determine tasking. Interpret goal/assignment into objects and actions of interest "Search buildings in area A to identify threats"	Plus: Ability to report findings in the perspective that best supports own troops' performance (ego or exo-centric, grid reference, "above to your 3 o'clock", etc.), knowing what the human wants to know and the different ways that information might be captured by available sensors	• Pre-existing knowledge base of objects and their properties (e.g., meaning of threat in this context) • Pre-existing knowledge base of actions that support goal (where to look/listen for threats in this context—windows of building, roof) • Pre-existing knowledge base of priorities by context (e.g., dwell time, search patterns)

[a] Definition from National Research Council (2002)

Table 3   Range of intent complexity (moves from taskings to reasoning about higher-level goal or mission or rules of behavior) (Continued)

Description	Variability	Preconditions
5. Adaptation from goal/assignment. Change or adapt mission/assignment based on changing situation	Plus: Ability to identify the need to break from current task, and complete unassigned, higher-priority task (recover own force causality, track high-value target ID'd, immediate danger) Plus: Adapt to phases of a normal mission progression (e.g., covertly scouting ahead to identify threats before unit reaches hostile location; then identifying clear avenues to move as enter area or responding to locations of sniper fire; or monitoring squirters[a] during building breach) Plus: Can reason on progress to assigned mission objectives and identify indicators of the failure of, or flaws in, the command concept/mission, and communicate information that could change or cancel the concept or to enact a contingency plan (e.g., enemy force much larger than planned, ambush or trap, hostages not present)	• Knowledge of priority hierarchy of tasks for the mission context • Knowledge of flow/tempo/stages of mission and indicators that a stage or phase has been reached (e.g., troops move into city, breach building, stop moving) • Knowledge of planned mission and end state (e.g., 5-paragraph orders) and related properties • The robot will need to understand key environmental aspects of the battlespace such as terrain and decisive actions/events/points—perhaps battle damage assessments or determination when key entities (e.g., enemies, clear path) are not where anticipated, invalidating plan.
6. Multihuman context. Robot can take input and provide information from various humans to support mission completion	Plus: Ability to support various teammates selectively with information or actions.	• Recognizing valid inputs on mission status/intent from multiple humans (sources) • Knowledge of the team composition and roles (doctrine) • Knowledge of which inputs to prioritize over others.

[a] "Squirter" is a person, assumed to be an enemy, running away from a military attack.

Table 3   Range of intent complexity (moves from taskings to reasoning about higher-level goal or mission or rules of behavior) (Continued)

Description	Variability	Preconditions
7. No consistent human input. Robot provided mission-level task/role based on endstate or objectives and can be left on autonomous duty to reason which tasks from a set of possible tasks it could best perform to support/ accomplish that mission given the current circumstances	Plus: Additional reasoning to determine trends and indicators (e.g., identify patterns of life or atmospherics—indicators of stability, pro/anti-US sentiment [including eavesdropping], or baseline movement patterns, meeting places, traffic, timing of activities, coherent change detection) Plus: If authorized to engage with lethal force the requirements will increase tremendously to determine priority of target, potential for collateral damage, priority of capture or kill, "hearts and minds", and the like. But this situation is not foreseen for the near-term future.	• Robot needs flexible mission tasking similar to CI (goals, endstate, mission tasks) • Robot needs model to help determine risk and rewards for various actions • Robot needs the ability to assess the situation to determine self-preservation and protection behavior

3.4 Assumptions

These are the assumptions that scoped and constrained the development or generalizability of list-intent behavior types presented in Sections 3.2 and 3.3.

Assumption 1. The focus is on spoken dialogue commands from the human and the natural language processing of taskings by the agent, but this does not rule out commands or "plays" given to the robot by other input methods such as via commands on a GUI. There must be other inputs, such as the pre-existing robot knowledge (knowledge-base), for almost any action to occur.

Assumption 2. To really test "intent", meaning the robot can adapt or act beyond the given task information provided, there needs to be ambiguity and flexibility, which is really only seen in types 4 and beyond.

Assumption 3. The reasoning does not have to be embedded on the actual robot platform; it is possible that there could also be an external computer-based reasoning component.

Assumption 4. The robot will not be continuously controlled remotely by the human and therefore must operate at least semi-autonomously.

Assumption 5. The focus here is on military contexts to constrain the wider field of possible human–robot interactions.

Assumption 6. These are phrased in terms of the human understanding what the robot will do and how to task it, meaning how much freedom of action and expected initiative the robot will exhibit in order to predict behavior and understand when the robot is supporting the intent and also when there is a disconnect. These are not phrased in terms of the most effective way to program the robot or reasoner.

Assumption 7. If robots are allowed the autonomous (not confirmed by humans) ability to engage targets (apply lethal force) a much wider set of rules and reasoning of intent will be required to evaluate priority of target, potential for collateral damage, priority of capture or kill, "hearts and minds", and the like. But this situation is not foreseen for the near-term future.

Assumption 8. It is possible for the robot/reasoner to clarify requests with the human when action is not clear but this was not listed for each type of intent behavior. The robot would still have to identify the need to ask/clarify and what needs to be clarified.

Assumption 9. The precondition knowledge items required are considered to build off each other as the reader continues through the list and therefore the items are not repeated for each type.

4.   Review and Conclusions

There is a consistent push for changing the ratio of multiple humans controlling one robot, or unmanned asset, to one human in charge of multiple robots or unmanned assets for military and other operations (Bray-Miners et al. 2012) even to the point of human and swarm interaction (Arquilla and Ronfeldt 2000; Crandall et al. 2017). It will not be possible for one human to carry out detailed control and monitoring of multiple assets (Chen and Barnes 2014). Therefore, concepts such as intent-based C2 relationships between human and robot teammates may prove essential for supporting these future paradigms.

The literature was reviewed and combined with input from practitioners to develop a range of intent-behavior types that could be realized in the military context, focused on those developed using natural language processing. It is envisioned that this work can provide a springboard for those looking to develop intent-based HRI for the military context. One goal is to start further discussion in the community as a way to refine and improve the framework. A second goal is to allow research and development personnel to look at the possible types of intent behaviors and ask what their systems are, and could be, capable of supporting and consider how to best map those capabilities on to the intended military scenarios and contexts. This will support asking the questions of what type of intent-based behavior is needed for a particular application context and what gaps do research and development efforts need to close to achieve that. The framework can also support a staged development plan, with researchers asking themselves what can be provided in 1 year, 3 years, or more to provide realistic expectations to both developers and end users.

As noted, this is a starting point. Further research is required to build the applications to test each of these types and really hammer out what is possible, what the limitations are, and most importantly what are the best ways to structure and support that HRI to provide the most effective, efficient, and satisfactory human–robot collaboration. One gap in particular is that this report takes a very human, operator-centric view on the problem. Additional input from the computer scientists and language and robotics experts is required as the development dialogue proceeds.

5. References

Arquilla J, Ronfeldt D. Swarming and the future of conflict. Santa Monica (CA): Rand Corporation; 2000.

Barnes MJ, Chen JYC, Hill S. Humans and autonomy: implications of shared decision-making for military operations. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2016 Dec. Report No.: ARL-TR-7919.

Bray-Miners J, Ste-Croix C, Morton A. Human–robot interaction literature review. Report prepared for the Canadian Department of National Defence for PWGSC; 2012, Toronto, Ontario. Contract No.: W7711-088136/001/TOR.

Bryant DJ, Webb RDG, Matthews ML, Hausdorf P. Common intent: a review of the literature. Technical report for the Canadian Department of National Defence for PWGSC; 2001, North York, Ontario. Contract No.: W7711-7-7404/01-SV. DCIEM No. CR-2001-041

Builder CH, Bankes SC, Nordin R. Command concepts: a theory derived from the practice of command and control. Santa Monica (CA): Rand Corporation; 1999.

Cannon-Bowers JA, Salas E, Converse SA. Cognitive psychology and team training: training shared mental models and complex systems. Hum Fact Soc Bul. 1990;33:1–4.

Chen J, Lakhmani S, Stowers K, Selkowitz A, Wright J, Barnes M. Situation awareness-based agent transparency and human-autonomy teaming effectiveness. Theor Iss Ergo Sci. Forthcoming.

Chen JYC, Barnes MJ, Wright JL, Stowers K, Lakhmani SG. Situation awareness-based agent transparency for human-autonomy teaming effectiveness. Presented at the International Society for Optical Engineering Defense and Commercial Sensing Conference; 2017 Apr 9–13; Anaheim (CA).

Chen JYC, Barnes MJ. Human-agent teaming for multirobot control: A review of the human factors issues. IEEE Trans Hum Mac Sys. 2014;44(1).

Clark HH, Brennan SE. Grounding in communication. In: Resnick LB, Levine J, Behrend SD, editors. Perspectives on socially shared cognition. Washington (DC): American Psychological Association; 1991. p. 127–149.

Clark HH, Schreuder R, Buttrick S. Common ground and the understanding of demonstrative reference. J Verb Learn Verb Behav. 1983;22(2):245–258.

Crandall J, Anderson N, Ashcraft C, Grosh J, Henderson J, McClellan J, Neupane A, Goodrich M. Human-swarm interaction as a shared control: achieving flexible fault-tolerant systems. HCII. Forthcoming 2017.

Dempsey R, Chavous JM. Commander's intent and concept of operations. Mil Rev. 2013 Nov–Dec:58–66.

[DOA] Headquarters, Department of the Army. MI gunnery for the military intelligence company of the brigade engineer battalion 1.0. Washington (DC): Headquarters, Department of the Army; 2016. Training Circular No.: 2-19.400.

[DOA] US Department of the Army. ADRP 5-0: The operations process. US Department of the Army; 2012 [accessed 2017 Feb 7]. https://armypubs.us.army.mil/doctrine/index.html.

[DOD] Headquarters, Department of Defense. Counterinsurgency (COIN) intelligence, surveillance, and reconnaissance (ISR) operations. Report of the Defense Science Board Task Force on Defense Intelligence. Washington (DC): Headquarters, Department of Defense; 2011.

Entin E. Optimized command and control architectures for improved process and performance. Arlington (VA): Office of Naval Research (US); 2005. Technical report.

Goodrich MA, Schultz AC. Human-robot interaction: a survey. Found Trends Hum Comp Int. 2007;1(3):203–275.

Grice HP. Speech acts: syntax and semantics. Vol. 3. In: Cole P, Morgan JL, editors. Logic and conversation. New York (NY): Academic Press; 1975. p. 41–58.

Gustavsson PM, Hieb M, Eriksson P, More P, Niklasson L. Machine interpretable representation of commander's intent. Presented at the 13th International Command and Control Research and Technology Symposium; 2008. http://dodccrp.org/events/13th_iccrts_2008/CD/Launch_CD.html.

Intelligence Sergeant, Fort Huachuca, Arizona. Personal communication, 2017 Apr 21.

Jurafsky D, Martin JH. Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition. In: Jurafsky D, Martin JH, editors. Chapter 19: Dialogue and conversational agents. Upper Saddle River (NJ): Prentice Hall; 2009.

Klein G. Characteristics of Commander's Intent Statements. Proceedings of the 1993 BRG Symposium on Command and Control Research; 1993.

Koschmann T, LeBaron CD. Reconsidering common ground: Examining Clark's contribution theory in the OR. In: Kuutti K, Karsten EH, Fitzpatrick G, Dourish P, Schmiddt K, editors. Proceedings of the Eighth European Conference on Computer-Supported Collaborative Work; 2003 Sep 14–18; Helsinki, Finland.

Levin DT, Killingsworth SS, Saylor MM. Concepts about the capabilities of computers and robots: a test of the scope of adults' theory of mind. Presented at the 2008 Human–Robot Interaction Conference; 2008 Mar 12–15, Amsterdam, Netherlands.

Marks MA, Mathieu JE, Zaccaro SJ. A temporally based framework and taxonomy of team processes. AMR. 2001;26(3):356−376.

McCann C, Pigeau R, English A. Using the command and control framework to analyse command challenges. Toronto, Ontario: Defence R&D Canada; 2002.

Merriam-Webster. Intent. Merriam-Webster Online Learner's Dictionary [accessed 2017 Aug 10]. http://www.learnersdictionary.com/definition/intent.

Murphy PJ. Forays into command intent: assessing the components of intent statements and developing a methodology for measuring shared intent. Edinburgh (South Australia): DSTO Systems Sciences Laboratory; 2002.

National Research Council. Technology development for Army unmanned ground vehicles: report by the Committee on Army Unmanned Ground Vehicle Technology. Washington (DC): National Academies Press; 2002.

Neely DS. Network-centric commander's intent: the key to network-centric warfare command and control. Newport (RI): Naval War College. 2003 May.

Oron-Gilad T. Scalable interfaces for operator control units; a common display to conduct MOUT operations with multiple video feed sources. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015. Final Research Report (Task 5).

Rosen MA, Bedwell WL, Wildman JL, Fritzsche BA, Salas S, Burke CS. Managing adaptive performance in teams: guiding principles and behavioral markers for measurement. HRMR. 2011;21:107–122.

Shattuck L. Communication of intent in distributed supervisory control systems [unpublished dissertation]. Ohio State University; 1995.

Shattuck L. Communicating intent and imparting presence. Mil Rev. 2000 Mar–Apr:66–72.

Spiker VA, Holder EH, Walls WF, Campsey WM, Bruce PF. SamePage: development of a team training tool to promote shared understanding. Fort Leavenworth (KS): Army Research Institute for the Behavioral and Social Sciences; 2007. Research Note 2007-05.

Sutton RS, Barto AG. Reinforcement learning: an introduction. Cambridge (MA): MIT Press; 1998.

[USMC] US Marine Corps. Combat Orders Foundations B2B0287. Camp Barrett (VA); Marine Corps Training Command; 2015. Student handout.

Wang N, Pynadath D, Hill S. Trust calibration within a human-robot team: comparing automatically generated explanations. Proceedings of the 11th ACM/IEEE International Conference on Human Robot Interaction; 2016 Mar 7–10; New Zealand. p. 109–116.

Winner JL, Freeman JT, Cooke NJ, Goodwin GF. A metric for the shared interpretation of commander's intent. Proceedings of the Human Factors and Ergonomics Society 51st Annual Meeting; 2007.

Appendix A. Review of the Findings by the Director's Strategic Initiative (DSI) Team

---

Intent has to be examined from a practitioner view outside of the literature as well. To support this, the meaning of intent was gathered from the perspectives of various members of the DSI team, representing various backgrounds and perspectives. Each team member was provided with a slide set presenting a summary of the literature review and open questions and asked to define what intent meant for their efforts. The slides are included in Appendix B. Answers are provided as follows with the respondent identified by their respective branch within the US Army Research Laboratory.

A.1 Responses

A.1.1 Response 1: Computational and Information Sciences Directorate (CISD)

Hi Eric,

The aspect of intent that I believe is relevant to Task 3 (adaptive robot behaviors) is information about "how low-level actions are to be executed by a robot." For intent-driven visual search that I'm working on, this includes:

* Time constraints (e.g., How long? How often? How many?)

* Spatial constraints (e.g., Where is most important? How close or far?)

* Types of objects of interest

* Required accuracy

* Importance relative to other simultaneous actions (urgency)

A.1.2 Response 2: Sensors and Electron Devices Directorate (SEDD):

Eric, I second the Computer and Informational Sciences Division email (response 1 above) on what robots "care" about when it comes to intent. Your slides address this well on slide 2, where fixed tasks become a multitude of tasks with changeable parameters addressed as "freedom of action" and "expected initiative". A fixed task (e.g. go to a waypoint) can have a changeable parameter representing freedom of action (speed) that addresses an expected initiative (urgency). I only use this as an example. I'm not sure speed/urgency is the best "scalable" parameter to use for this DSI, but it represents an aspect of mobility that can go from being specified to being "servoable" based on perceived urgency.

Please try to assume the robots will NOT be employing learning techniques for the scope of this DSI (slide 4 of your attached deck). Typical (default) behaviors should be "hard coded". Scalable capabilities should be isolated during experiments to definitively show some aspect of intent caused a change to how the robot behaves,

whether this is in mobility or in perception. Think of this as a single intent input/single output behavior type setup for the purposes of getting this going.

Slide 6 (scenario requirements to test intent): Each of these bullet points is probably in the scope of this project, but should probably not be done at the vehicle/robot level. These probably need to be determined at the reasoner/planning area that's aggregating the 3-D world model.

A.1.3 Response 3: Vehicle Technologies Directorate (VTD)

Eric,

For the hybrid vehicle part, and task 3 generally I think. I'm expecting the intent to be parsed and rather concrete to be provided through the tactical behavior specification (TBS). Key things which need to be conveyed in the TBS include information about location/targets and urgency.

A.1.4 Response 4: Computational and Information Sciences Directorate (CISD)

Hi Eric,

I understand "intent" to mean "what the human wants the robot to do."
With respect to my piece of the DSI, one thing it means is "what the human is referring to in space." With respect to building predictive gaze models, another thing it means is "what the human is currently trying to do."

A.1.5 Response 5: Computational and Information Sciences Directorate (CISD)

For Task 2, I am considering any knowledge of context that alters how a natural language command is turned into a tactical behavior specification by the reasoner as "understanding intent." This includes prior statements of mission purpose, knowledge of the surrounding area, and some understanding of universal human goals.

A.1.6 Response 6: Human Research and Engineering Directorate (HRED)

The question is how much inferencing does the robot do. If the robot is told that the task is urgent and it increases its speed- is that inferencing or simply "hard coding". As your slides indicate- inferencing is a little more complex 1. Type of mission 2. Terrain 3. Urgency and inferencing a. are soldiers in danger b. is the mission usually urgent, etc. From our point of view, it may even be more complicated - what are we evaluating-- doubt we would investigate a robot with intent vs. a robot without intent -- even different levels of intent may not fit into this paradigm.

Intent - I think would mean the robot is given general mission and it infers based on its world model (the "reasoner); however it seems at this juncture that the "reasoner" will be pre-set to solve specific questions that are the ones being tested. This might be tricky for real world as you cannot pre-set your problems and hard code a solution -unless it is a complicated rule-based inferencing engine with many rules (hundreds -thousands with all the problems that -that involves-- obviously not).

A1.7 Response 7: Computational and Information Sciences Directorate (CISD) (with responses from the author of this report following).

A few thoughts and questions:

1. Is Intent the same as Intention?

2. What is the relationship between intent and desired goal?

3. If you have an intention, do you also have in mind a commitment to try to make that intention happen?

OK, now some comments, not questions:

1. It's important to make distinctions between intent that involves actions vs effects. Actions are clear - intention to walk, move, jump, etc. Effects are less clear, like to make someone feel guilty, to control a conversation. These are related to the Austin perlocutionary issues we talked about on Monday.

2. There are intentions that require multiple actors (e.g negotiating a price) and those that are one-sided. For the DSI, these are important to differentiate since the first requires the human and machine to "work together" to achieve a goal, and the second doesn't.

So to answer your question, Eric:

Intent is a commitment on the part of an agent (human or robot) to arrive at a desired state or achieve a desired result; intent could involve a single actor or multiple actors; intent requires common definitions of contextual relevant concept and elements.

A.1.7.1 Author's Response to Response 7:

Hi, thanks for the questions, thoughts and comments. You bring up very good questions and based on the feedback I have gotten so far the teammates (at least for robots) are thinking more on the lines of actions than effects, specifically commanded behaviors for the robots and the intention is largely solving spatial ambiguity and parameters of that task (e.g., speed, priority, etc.), rather than goal

(mission)-based ambiguity. So far only one response mentioned purpose, knowledge or goals which is of course the one you have included from CISD (response 5 above).

To attempt very novice answers to the questions (embedded in CAPS)

Subject: RE: Your understanding of intent for this DSI (UNCLASSIFIED)

A few thoughts and questions:

1. Is Intent the same as Intention?

FROM MY BACKGROUND UNDERSTANDING AND FOR THIS PROJECT'S USE I WOULD SAY NO.

THE LINK FROM INTENTIONS TO BEHAVIOR IS PRETTY WEAK IN THE RESEARCH AND IN THIS CASE ANY INTENT IS MEANT TO BE RESOLVED BY THE ROBOT IN AN ACTIVE BEHAVING CONTEXT. THEREFORE MY HYPOTHESIS IS THAT INTENT IS MEANT MORE AS A COMMAND/TASKING, PERHAPS WITH ASSOCIATED BUT UN-DECLARED GOALS, TASKS AND CONSTRAINTS

2. What is the relationship between intent and desired goal?

I THINK THIS IS THE NUT WE ARE TRYING TO CRACK, AS IN SHOULD THIS INTENT BE CONVEYED BETWEEN HUMAN AND ROBOT (OR REASONER AS MIDDLE MAN) AS A LARGER GOAL (STABILIZE AREA), A MISSION (PUSH ENEMY OUT OF THIS TOWN), A DETAILED TASK (SEARCH FOR THREATS IN THIS SPECIFIC AREA/BUILDING AND REPORT LOCATION) OR EVEN MORE BASIC (GO TO LOC X AND LOOK) AND FOR EACH OF THESE HOW IS SUPPLEMENTAL KNOWLEDGE GOING TO BE INTRODUCED TO ENSURE UNDERSTANDING AND HOW WILL PROGRESS BE MONITORED TO MAINTAIN THAT UNDERSTANDING? ARE WE ALSO CONSIDERING GOALS AND EFFECTS TO MEAN THE SAME THING? A COMMANDER CAN HAVE A LIST OF DESIRED EFFECTS TO ACHIEVE AND TASKS RELATED AS SUPPORTING THOSE OR NOT. IF WE START TO WORK AT THE HIGHER LEVELS OF ABSTRACTION (WHICH WOULD BE ENGAGING) THEN WE WILL HAVE TO REALLY PUT IN SOME TIME TO THINK ABOUT AND BUILD LEVELS OF GOALS THAT FIT THE CONTEXT OF INTENT (WHAT SUB-GOALS, TASKS, ETC. FIT THAT CONTEXT AND PERHAPS ADD PRIORITIES). FOR INSTANCE IF THE ROBOT/REASONER KNOWS WE ARE ON A SPECIFIC KIND OF MISSION WITH SPECIFIC GOALS AND SUB-GOALS AND BEHAVIORS THAT TYPICALLY SUPPORT THIS CAN IT BE ASSIGNED A

TASK AND ADAPT BEHAVIOR AS A SCENARIO PROGRESSES REMAINING WITHIN THAT OVERALL GOAL OR INTENT? THE EXAMPLE FROM ANOTHER PROJECT WAS THE DIFFERENCE BETWEEN GIVING THE GOAL TO GO TO THE BACK DOOR TO WATCH FOR PEOPLE OR INSTILLING SOME SENSE OF THE HIGHER TASK OF REPORTING ANY ENEMY MOVEMENT OUT OF THE BACK OF THE BUILDING (THE ROBOT WAS SLOW AND ENEMY RAN PAST IT ON THE SIDE OF THE HOUSE AND IT DID NOT REALIZE THAT WAS IMPORTANT). HOPEFULLY WE MOVE TOWARDS THE LATTER BUT HOW MUCH OF THAT REALLY COMES DIRECTLY FROM NLP WITHOUT BUILDING UP A COMPLICATED CONTEXT-BASED REASONER I CANNOT SAY.

3. If you have an intention, do you also have in mind a commitment to try to make that intention happen?

THIS IS WHAT I NOTED IN NUMBER 1. I WOULD STAY AWAY FROM INTENTION AS A KEYWORD AS IN THE BEHAVIORAL SCIENCES THE LINK BETWEEN INTENTION AND ACTION IS WEAK AND FOR ROBOTS THAT SHOULD NOT REALLY BE THE CASE. IF THERE IS AN INTENT COMING FROM A COMMANDER/OPERATOR IT IS LIKELY NOT VENTING OR DAYDREAMING OR ANYTHING SHORT OF A CALL TO ACTION WHICH BREAKS FROM THE HUMAN RESEARCH.

OK, now some comments, not questions:

1. It's important to make distinctions between intent that involves actions vs effects. Actions are clear - intention to walk, move, jump, etc. Effects are less clear, like to make someone feel guilty, to control a conversation. These are related to the Austin perlocutionary issues we talked about on Monday.

2. There are intentions that require multiple actors (e.g negotiating a price) and those that are one-sided. For the DSI, these are important to differentiate since the first requires the human and machine to "work together" to achieve a goal, and the second doesn't.

DO WE CONSIDER THE ROBOT/REASONER TO HAVE ITS OWN INTENT OR JUST A PROGRAMMED SEQUENCE? MY IMPRESSION SO FAR IS THAT THE INTENT IS MOSTLY UNI-DIRECTIONAL HUMAN TO ROBOT WHO THEN HAS TO UNDERSTAND AND ACT UNDER IT. THERE COULD BE MORE THAN ONE HUMAN INPUTTING INTENT BUT I DO NOT THINK WE ARE GOING TO COMPLICATE IT LIKE THAT EXCEPT FOR PRE-SET NON-VERBAL INPUTS INTO THE REASONER. THE COORDINATION COMPONENT SHOULD BE THERE BUT NOT LIKELY AT THE COGNITIVE

LEVEL OF HUMAN COORDINATION BUT RATHER THE IF SEE A DO B AS IT FITS MY ORDERS (INTENT), FOR EXAMPLE, THE HUMAN TEAMMATES HAVE MOVED TO THIS STAGE OF MISSION, SO ROBOT SWITCHES ITS TASK BEHAVIOR TO BEST SUPPORT THAT STAGE (FROM SEARCH EVERYWHERE TO OVERSIGHT OF SPECIFIC PATH OF APPROACH)

So to answer your question, Eric:

Intent is a commitment on the part of an agent (human or robot) to arrive at a desired state or achieve a desired result; intent could involve a single actor or multiple actors; intent requires common definitions of contextual relevant concept and elements.

I THINK THAT DEFINITION MOSTLY FITS BUT MIGHT MISS THE ONE-WAY DELIVERY COMPONENT OF INTENT WE MIGHT BE WORKING UNDER, AS THERE IS NO REAL COMMITMENT FROM A ROBOT. IN MY UNDERSTANDING IT IS MORE OF A TASKING THAT IS FOLLOWED AS BEST THE ROBOT CAN REASON TO DO SO RATHER THAN A COMMON CAUSE. THE INTENT MIGHT BE MORE THE DEFINING AND CONVEYING OF THE END STATE, AND TO SOME TO BE SEEN DEGREE, HOW TO SPECIFICALLY GET THERE. I THINK I PUT OUT MORE QUESTIONS THAN ANSWERS AND QUALIFY ALL COMMENTS AS BEING FROM JUST ONE PERSPECTIVE THAT MAY OR MAY NOT FIT THE LARGER "INTENT" OF THE PROJECT. BUT THAT IS WHAT I AM PUTTING THIS STUFF OUT THERE IN HOPES OF FIGURING OUT.

A.2 Summary

To summarize, what can be seen in these different perspectives is that each stakeholder can have a different view of what intent means, as well as the extent of adaptability under the umbrella of intent and this is colored by his or her background and project tasking. The challenge from this point was to develop the range of intent-based behaviors based on the literature as well as input from the team to provide guidance for current and future development efforts. This was the focus of Section 3 above.

INTENTIONALLY LEFT BLANK.

Appendix B. Briefing Provided to Director's Strategic Initiative (DSI) Team to Collect Intent Definitions

This appendix appears in its original form, without editorial change.

Approved for public release; distribution is unlimited.

47

U.S. ARMY RDECOM

HRI Intent for NLP Context

ARL

What can we learn from human-human interaction concerning intent and what applies to human-robot?

- Lots of different views and theories in the literature
- Most do not apply "as-is" to our human-robot context
- Need to have a team discussion to help define scope
  - Potentially-relevant concepts for intent-inferencing
  - The context we are focusing on
  - Initial intent: Explicit vs implicit information
  - Other methods of building or maintaining shared intent
  - Scenario needs to truly test intent

U.S. ARMY RDECOM

HRI Intent for NLP Context

ARL

Potentially-Relevant Concepts for Intent-Inferencing

- Command(er's) Intent (CI): purpose, method, end-state w/ <u>no instructions</u> to allow freedom (2 levels)
  - Auftragtaktic (German version of CI)
  - Commander's Vision
  - Shattuck's Presence
- Mission-type Orders: larger in scope than CI to include task and purpose but not TTPs
  - 5-paragraph order
  - Builder's Command Concept (Vision)
- Common Intent (Pigeau & McCann): wider breadth and includes explicit/implicit-- expectations, beliefs, and values borne out of personal, military, and cultural experience
  - Sensemaking (Brehmer)
  - Mission Awareness
- Hierarchical Goal Representations
  - By types of operations
  - Effects-based operations
- Shared Understanding (focused on items of interest)
  - SamePage Model of factors
  - Shared Mental Models: task, team, equipment/resources
  - Common Ground and grounding research
- Address **Freedom of Action** and **Expected Initiative** for Robot teammates
- Could also focus more on teamwork and taskwork literature

U.S. ARMY RDECOM HRI Intent for NLP Context **ARL**

How our Context of interest impacts focus

- Planning Context: more strategic and preparatory and where CI and orders usually come in
- Mission Execution: More dynamic with tactical coordination and response interactions and understanding
- Is focus on understanding shared intent by individual (commander), team, or robot?
- One time, one mission, or longitudinal learning by team and robot (machine learning?)
- Level of command: squad or platoon, right?

U.S. ARMY RDECOM HRI Intent for NLP Context **ARL**

Initial Intent: Explicit vs Implicit Knowledge

- How will each be communicated to robot: pre-programmed, based on briefing-NLP, machine learning, LSA, etc.?
- Explicit- what to include:
    - Mission orders and plans (higher level too), verbal commands, activating specific programs, etc.
    - What about other team member roles, specialties and knowledge (coordination training). Should/could be specific to mission task context.
    - Gaze direction included but how about tone of voice and other cues
- Implicit- what to include:
    - Are beliefs, values, etc. relevant to robots and which ones and how to portray (rules, constraints)?
    - Should we code various ROEs or other if-then rules and priorities (e.g., human life, targeting lists and priorities, who needs what information, whose input/commands to respond to, etc.)
    - Could the robot learn from past vignettes, simulators, lessons learned, hero stories, field manuals, doctrine, etc.?
- Are either of these knowledge types context dependent?
    - If the robot learns from experience should this be applied to other situations, robots, etc.?
- Battle Management Language and similar means for coding this knowledge?

Methods for Building and Maintaining Shared Intent:

- Leverage common cues/methods in human-human understanding
  - Either robot understands or uses cues/methods
  - Examples: backbrief, discussion, clarification requests, rehearsals, grounding methods--uh-huhs, closure, etc.
- When to use ego or exocentric references (context-dependent, also ground vs aerial)
  - Annotating maps/imagery is one means often used to build understanding
- Reporting on key events that fit intent (i.e., objectives achieved)
- Asking questions when unclear on meaning or how to proceed
  - Shattuck's 6 categories of verbalizations
- Providing information to right teammate at right time without request (anticipation ratios)
  - Might need to be on common communication channel but with intended recipient specified
- Adjusting communication or reporting strategies based on mission variables (phase, workload, need for silence, other adaptation conditions)
- Communication medium can have impact on cues and feedback so should be considered (voice, message, face-to-face, etc.)

Scenario requirements to test intent:

- Need responses to unpredicted/unanticipated events
  - If not we are just following a plan which does not test adaptation under intent
  - Example from CTA of robot tasked to guard door not reporting enemy fleeing early because did not understand larger task
- Need decision points where the robot has to choose, report for guidance, or make collaborative decisions
  - Which suspect to follow (prioritized target sets or mission goals)
  - Friendly or neutral identification
  - Proceed to next task/location

How to transition to human-robot teams?

So what does this mean for human-robot teams in our DSI context?

The Nation's Premier Laboratory for Land Forces

List of Symbols, Abbreviations, and Acronyms

ADRP	Army Doctrine Reference Publication
ARL	US Army Research Laboratory
C2	Command and Control
CE	Controlled English
CI	Commander's Intent
COA	courses of action
COP	common operating procedure
DSI	Director's Strategic Initiative
GUI	graphical user interface
HRI	human–robot interaction
NLP	Natural Language Processing
ROE	rules of engagement
SICI	Shared Interpretation of Commander's Intent
SU	shared understanding
SITREP	situation reports
TBS	tactical behavior specification
ToM	theory of mind

INTENTIONALLY LEFT BLANK.