Shared Understanding within Military Coalitions: A Definition and Review of Research Challenges

Paul R. Smart[1*], David Mott[2], Katia Sycara[3], Dave Braines[2], Michael Strub[4] and Nigel R. Shadbolt[1]

[1]*School of Electronics and Computing Science, University of Southampton, Southampton, SO17 1BJ, UK*
[2]*Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, Hampshire, SO21 2JN, UK*
[3]*Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, 15213-3890, USA*
[4]*Army Research Laboratory, Alexandria, Virginia, 22310-3225, USA*

[*]*Corresponding author: ps02v@ecs.soton.ac.uk*

Abstract—*Shared understanding is commonly seen as essential to the success of coalition operations. Anecdotal reports suggest that shared understanding enables coalition forces to coordinate their efforts in respect of mission goals, and shortfalls in shared understanding are frequently cited as the reason for poor coalition performance. In spite of this consensus regarding the importance of shared understanding, however, there are very few empirical studies that attempt to explore shared understanding in a military coalition context. This paper attempts to support future research efforts into shared understanding by proposing a specific definition for shared understanding and identifying a number of research challenges. Shared understanding is defined as the ability of multiple agents to exploit common bodies of causal knowledge for the purposes of accomplishing common (or shared) goals. This definition implies that agents possessing shared understanding will be capable of coordinating their respective behaviours in order to ensure the efficient realization of cognitive and behavioural objectives. We also identify a number of areas for future research into shared understanding. These include the factors that affect shared understanding, the effect of shared understanding on coalition performance, and the development of techniques to reliably measure and assess understanding in coalition environments.*

1. INTRODUCTION

Shared understanding emerges as a construct of considerable significance in discussions about coalition operations. Anecdotal reports suggest that shared understanding enables coalition forces to coordinate their efforts in respect of mission goals, and shortfalls in shared understanding are frequently cited as the reason why coalition forces fail to realize military objectives (Maj. Edward Gentle, personal communication). Despite the apparent importance of shared understanding to coalition operations, there is little, if any, agreement as to what the term 'shared understanding' actually means. Few studies have sought to examine shared understanding in coalition contexts, and there is little hard scientific evidence regarding the factors that promote or undermine shared understanding in coalition operations. In order to address this shortcoming, the current paper attempts to review some general issues and challenges associated with the analysis of shared understanding in military coalition contexts. In particular, the current paper attempts to address the following specific questions:

1) What is meant by the terms 'understanding' and 'shared understanding'?
2) What is the relationship between shared understanding and ostensibly similar constructs, such as shared situation awareness (SSA) and shared mental models (SMMs)?
3) Do the notions of understanding and shared understanding apply to synthetic agents? Is it possible to have shared understanding in hybrid agent teams?
4) What forms of shared understanding are encountered in collaborative task contexts?
5) What are the specific research challenges that need to be addressed by future research efforts?

In line with these questions, Sections 2 and 3 propose definitions for the notion of understanding and shared understanding, respectively. A key aim of these sections is to distinguish notions of understanding from related constructs that have been investigated in the literature, e.g. situation awareness (SA). Section 4 tackles the notion of machine understanding. It maintains that our definition of understanding is consistent with the possibility that shared understanding can be a feature of hybrid agent teams. Section 5 discusses a number of ways in which understanding may be shared or distributed across individuals. It also introduces the notion of distributed understanding, which is the idea that the realizing mechanisms for understanding may, at times, be distributed across a variety of processing elements and material resources. Finally, Section 6 suggests a number of issues to be addressed by future research into shared understanding.

2. WHAT IS UNDERSTANDING?

In order to derive a definition of shared understanding, it is important to understand what we mean by the term 'understanding'. The notion of understanding that we will countenance here owes much to Wittgenstein's [1] notion of understanding as a kind of ability. Understanding for Wittgenstein was akin to the possession of a technique, or skill, hence the origin of the slogan that understanding is

"meaning in use" [1]. Our definition of understanding places a similar emphasis on ability:

> *Understanding is an ability to exploit bodies of causal knowledge (i.e. knowledge about the antecedents and consequents of particular phenomena) for the purpose of accomplishing cognitive and behavioural goals.*[1]

To understand something (e.g. a situation, action, linguistic utterance, and so on) is thus to be able to do things[2] that depend on a knowledge of causal relationships about how that something came to be, and what effects that something is likely to have. In the human case, we suggest that understanding comprises an ability to, *inter alia*, establish veridical[3] expectations and explanations about a variety of phenomena (events, situations, actions, system states and so forth). Such expectations and explanations are constitutive of understanding, we suggest, because they represent goal-relevant behaviours that depend on the exploitation of bodies of (implicit or explicit[4]) causal knowledge.

Our characterization of understanding as an ability to form expectations and explanations aligns itself with a body of recent work concerning mental models [2-4]. According to Rouse and Morris [5] mental models are the "mechanisms whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future system states" (pg. 351). Mental models therefore appear to play a key role in *enabling* understanding: they support the generation of behaviours that warrant talk of an individual as understanding some aspect of a domain.

In addition to mental models, it is important to consider how notions of understanding relate to the concept of SA and SSA. According to Endsley [6], SA is "the perception of the elements in the environment within a volume of space and time, the comprehension of their meaning and the projection of their status in the near future" (pg. 36). Two elements of this definition seem immediately relevant to our notion of shared understanding: comprehension and projection. Comprehension implies that individuals who possess SA are capable of interpreting, combining and prioritizing information. Projection, on the other hand, implies that individuals are capable of formulating expectations about the occurrence of future events and system states. On the basis of Endsley's [6] characterization it would therefore seem that understanding is something that is subsumed by SA – that the possession of SA necessarily entails understanding. To our mind this does not seem entirely appropriate. It seems that one could be aware of situation-relevant information without necessarily understanding that information. For example, one could be aware of one or more items of information without necessarily adjusting one's response output in goal-appropriate ways (i.e. ways that ensure the efficient realization of behavioural and cognitive goals). SA, we suggest, might be better conceived of as the functional poise of situation-relevant information to influence thought and action throughout the course of online, situation-directed behaviour. This alternative conceptualization might achieve some degree of theoretical separation between the notions of SA and individual forms of understanding[5].

3. SHARED UNDERSTANDING

Based on the definition of understanding outlined above, we define shared understanding as follows:

> *Shared understanding is the ability of multiple agents to exploit common bodies of causal knowledge for the purpose of accomplishing common (or shared) goals.*

As with the notion of individual understanding, shared understanding entails a commitment to the idea that understanding is an ability; namely an ability to adaptively modify thought and action in ways that ensure the efficient realization of cognitive and behavioural goals. The claim that common bodies of causal knowledge are exploited for the purposes of accomplishing common goals implies (but does not necessarily entail) that the cognitive and behavioural responses of agents possessing shared

[1] The definition of understanding we propose is intended to serve as the basis for analyses of shared understanding in military coalition contexts (see Section 3). We make no claims about the broader applicability of the definition to other usage contexts. It may be that the definition can accommodate other notions of understanding (e.g. language understanding, situation understanding and the understanding of intentional action), but it may also be the case that there are multiple types of understanding each requiring a separate definition.

[2] Note that our emphasis on understanding as an ability does not entail a commitment to distinguishing between dispositional and occurrent 'forms' of understanding. To entertain this view (of multiple forms of understanding) would be to suggest that understanding (in a dispositional sense) is a disposition to understand something (in the occurrent sense) if a suitable occasion for behavioural (e.g. linguistic) expression should arise. This is prohibited on our account because to understand something (e.g. a situation) is to be able to express behaviours in a manner that is aligned with causally-relevant contingencies and goal states. Behavioural competence (broadly construed) is, on our view, constitutively relevant to understanding.

[3] Note that while our notion of understanding in the individual case depends on the veridicality of explanatory and predictive inferences, the notion of *shared* understanding (to be developed later) does not necessarily entail this commitment to veridicality. Two or more individuals could, theoretically, share an understanding even if that understanding was inaccurate or mistaken. The same cannot, it seems, be true of understanding in the individual case (an individual only understands something if their predictive inferences are veridical). This is one example where notions of understanding in the individual case do not appear to coincide with the notion of understanding in the shared case. That is, shared understanding does not appear to be the mere sharing of individual forms of understanding.

[4] We suggest that an *explicit* knowledge of causal contingencies, such as might be required in formulating explanations of system behaviour, is not a prerequisite of understanding. This is somewhat controversial because some might argue that explicit, linguistically-formulated explanations are a necessary feature of understanding. We see no principled reason to motivate this conclusion. Understanding, on our view, simply requires an ability to exploit causal knowledge in order to adaptively regulate response output in goal-appropriate ways. In some cases the behavioural outputs may be language-based, but non-linguistic forms of behavioural output seem equally valid.

[5] Ultimately, the issue of whether SA can be studied independently of understanding is irrelevant. We assume that the notion of understanding is more generic than the notion of SA and that its ontological integrity does not depend on whether SA necessarily entails understanding.

understanding will be highly similar (or at least functionally equivalent with regard to goal realization[6]). In fact, we might conclude that similarity of response output is one way to measure the extent (or overlap) in the understanding possessed by agents. Thus, two individuals who possess shared understanding will, at least in some cases, establish the same set of explanations and expectations given identical information about (e.g.) system states (all other things being equal). In the case of medical diagnosis, for example, we might conclude that two individuals have the same understanding if they are able to account for symptoms in the same way, and are additionally able to anticipate the same set of pathophysiological outcomes as a result of disease progression. In a coalition military context we might say that two commanders have the same (i.e. shared) understanding of a situation if they are able to anticipate the same effects of military actions, and are also able to cite the same reasons as to why particular military actions should be undertaken (e.g. to ensure the efficient realization of mission objectives). Clearly, the shared understanding that individuals possess (as measured by predictive and explanatory capabilities) will not be identical in most cases. In addition, the shared understanding between individuals will rarely, if ever, be complete. More likely, individuals will possess limited forms of shared understanding that are specific to some situation or task context (see Section 5 for more on this).

Although we have talked about shared understanding in terms of commonality of response outputs, it is important to be clear that we do not see commonality of response outputs as necessary for shared understanding. An alternative view of shared understanding, one that is commonly encountered in discussions regarding coalition operations, is that shared understanding entails an ability to coordinate the thoughts and action of multiple individuals so as to ensure the efficient realization of some common or shared goal. This view sees shared understanding as contributing to something like unity of effort, the notion that coalition force elements (perhaps from different command structures) are able to cooperate and coordinate effectively in order to realize common mission objectives. Following on from this view, we might be inclined to define shared understanding as something like the following:

Shared understanding is the ability of multiple agents to coordinate their behaviours with respect to each other in order to support the realization of common goals or objectives.

In this case, understanding does not imply commonality of response output, because each agent may have to undertake different actions in order to ensure that some common goal is accomplished. It is clear that when common goals are being pursued, agents may sometimes need to adapt their response outputs in different ways based on an awareness of what other agents are currently doing. This is particularly true in the case of military coalition operations, and it is therefore important that our definition of shared understanding should not exclude this kind of coordinative function. We suggest that definitions of shared understanding should not oblige us to accept commonality of response output as a necessary condition for shared understanding. Rather, we should see shared understanding as an ability to adaptively modify behaviours in ways that ensure collective goals are accomplished. In some situations, namely those in which collaboration is not required, then shared understanding will be indicated by common behaviours (see the above case of medical diagnosis). In other cases, shared understanding will be indicated by an ability to engage in different (agent-specific) responses that are nevertheless coordinated with respect to each other. What is common to these two cases, we argue, is the ability of agents to exploit bodies of causal knowledge in order to adapt thought and action in ways that ensure the attainment of common goals. This is precisely the target of our proposed definition of shared understanding. The definition does not limit us to commonality of response output, although it is likely that in most cases commonality of response output will be the easiest means by which to measure shared understanding.

As with understanding in the individual case, it is important to distinguish between shared understanding and ostensibly similar notions such as SMMs and SSA. SMMs [2, 3] are mental models that are possessed by multiple individuals. They are assumed to benefit team coordination and communication because they enable individuals to anticipate one another's information requirements and interpret events in similar or identical ways. Inasmuch as mental models provide a realizing mechanism for individual forms of understanding (see Section 2), it is possible that SMM may provide one means by which shared forms of understanding may be realized. Nothing in our definition of shared understanding, however, commits us to the idea that individuals *must* have similar or identical mental models in order to possess shared understanding[7]. In fact, we maintain that, at least in some cases, the actual details of the physical mechanisms that realize shared understanding are largely irrelevant in terms of our efforts to delineate the relationship between shared understanding and group interaction/coordination processes. In other words, it does not matter how individuals manage to realize the behaviours that warrant the ascription of understanding; what matters are the behavioural outputs themselves. Something like this

[6] The notion of functional equivalence is included here to account for the fact that two more response outcomes may be equivalent with regard to the realization of specific goals. In these cases, it does not matter which response is selected since neither is better or worse than the other. Differential response selection in these cases is no guide to the actual similarity of understanding possessed by agents, because agents could possess the same understanding and yet select different responses using arbitrary criteria. One potential way round this problem is to suggest that what is important is not so much similarity of response choice as similarity of response generation. Thus, while individuals may select different (functionally equivalent) responses based (e.g.) on personal preferences, if they possess shared understanding they will nevertheless generate similar sets of response alternatives from which a selection is made.

[7] This aligns itself, to some extent, with the notion of equifinality that is discussed in the SMM literature [see 7].

conclusion is also apparent in the SMM literature. Cannon-Bowers et al [3] thus argue that because the "function or benefit of shared mental models is that they lead to common expectations of the task and team, it is the expectations rather than the mental models themselves that must be shared".

Understanding the relationship between shared understanding and SSA is somewhat more difficult than is the case for shared understanding and SMMs. This is partly because the notion of SSA includes the same elements of comprehension and projection that proved so problematic in the case of individual SA [8-10] (see Section 2). As with individual SA, it may be necessary to discriminate between a shared awareness of situation-relevant information and the common (e.g.) predictive and retrodictive capabilities that constitute shared understanding.

4. MACHINE UNDERSTANDING

The notion of shared understanding that we developed in Section 3 implies that multiple individuals (or agents) possess similar (or identical) abilities when it comes to (e.g.) the prediction and explanation of (e.g.) system behaviour. This view does not entail a commitment to the idea that similar mechanisms need to undergird the manifestation of these abilities. Different agents could use very different mechanisms to generate explanations and predictions without necessarily undermining the possibility that they possess shared understanding. All that is required for shared understanding, in our view, is an ability to adapt response output with respect to bodies of causal knowledge in ways that lead to the realization of shared goals. We therefore embrace a functionalist view of shared understanding. We allow for the possibility that shared understanding is independent of the specific details associated with realizing mechanisms. Given these functionalist intuitions about shared understanding, we suggest that machine agents can possess understanding, and that they can share their understanding with other agents.

One of the major objections to the notion of machine understanding derives from the philosophical community. Searle's [11] famous case of the Chinese room tends to undermine the idea that computational operations defined across formally specified elements can yield the kind of understanding that characterizes biologically-circumscribed cognition. Searle's [11] thought experiment has been the subject of a lively debate about the tenability of formal accounts to yield genuinely intelligent behaviour; however, it is not our purpose, in this paper, to refute or rebut Searle's argument. Instead, we suggest that in concluding that no amount of formal syntactic manipulation could (ever) yield genuine understanding, some theorists may have relied too much on an overly simplistic (and poorly defined) notion of what understanding really means. The notion of understanding, as discussed by Searle [11], does not, we suggest, correspond to the notion of understanding that is presented here. And arguments against formal, information-theoretic formulations of machine understanding thus gain little or no leverage with respect to the present discussion. The logic of Searle's [11] argument seems to be that understanding consists in something more than an ability to engage in adaptive sequences of goal-relevant response output. This additional feature is, we suggest, perhaps something akin to the phenomenal experience of understanding – the conscious experiences we have whenever we 'feel' as though we understand something. Given this apparent emphasis on phenomenal experience, it is not surprising that formal systems (at least of the kind described by Searle) seem unlikely to possess any kind of understanding. The problem is that we do not accept the claim that conscious experiences are indeed constitutively relevant to understanding (at least the form of understanding that is being discussed here). There are two reasons for this[8]. Firstly, human subjects may feel that they understand something, even though they may be mistaken. Secondly, it seems that we may, on occasion, understand something even when the conscious feelings associated with understanding are lacking[9]. One (admittedly theoretical) example of this is Price's [12] account of why we find it so hard to accept that we have some understanding of consciousness in spite of the fact that we do seem to have access to a significant body of explanatorily-relevant information (consider the wealth of knowledge derived from neuroscientific and neuropsychological analyses). What Price [12] suggests is that our problems in understanding consciousness are not attributable to the ontological or metaphysical profile of phenomenal experience per se; rather, the problem is rooted in how the feelings of understanding are generated in the first place. The "warm glow of explanatory understanding", Price [12] suggests is the result of a kind of self-deception in which we fool ourselves into seeing effects as contained within their causes [see also 13]. The point, for present purposes, is simply that our conscious feelings regarding the depth of our understanding are generally poor guides as to our actual level of understanding in some domain. Inasmuch as Searle's [11] claims are based on the idea that phenomenal experience is a necessary part of understanding, we suggest that such claims cannot establish the case against machine understanding. Synthetic agents could, we suggest, engage in behaviours that warrant talk of them as legitimately understanding some state of affairs. We also suggest that inasmuch as these behaviours (e.g. predictions about future events) parallel those made by other agents (including human agents) then such agents (human and machine) should be considered as possessing shared understanding relative to each other.

[8] Potentially, a third reason exists. It is that while multiple individuals seem capable of expressing the same cognitive, behavioural and linguistic competencies (as is the case in our notion of shared understanding), it is less clear that they can share conscious experiences. If we conclude that 'feelings' are of constitutive relevance to understanding then the possibility that we can ever encounter genuine forms of shared understanding (in which the same conscious experience must be shared) seems problematic.

[9] There is, of course, an additional argument here. It is that, even in the human case, we may encounter forms of implicit understanding that do not necessarily entail conscious experiences.

5. SHARING UNDERSTANDING

As Cannon-Bowers and Salas [14] point out in the context of shared cognition, the notion of 'sharedness' can be viewed in multiple ways. One interpretation of 'shared' is that it denotes the common or joint possession of some resource (e.g. the sharing of a belief or item of equipment). An alternative view sees 'sharing' as implying the division of a resource between multiple recipients (e.g. sharing the workload or sharing a dessert). The former view of sharedness is clearly the one that is most relevant to the notion of shared understanding as discussed in this paper. On our view, shared understanding is the possession of a single resource (i.e. an ability), and that resource is not divided into parts and distributed across multiple agents. Perhaps there are some situations, however, in which it does make sense to talk about understanding as being shared in the distributed sense of the term. Perhaps, in such situations, it is collections of individuals that possess understanding, not individuals themselves.

In this section we attempt to present a variety of views of understanding, each of which differs with respect to the way in which the abilities that constitute understanding are distributed among a variety of processing elements and material resources.

5.1. Identical Understanding

One view of shared understanding emphasizes the need for agents to possess identical abilities. This form of shared understanding assumes that the understanding possessed by one or more agents is completely overlapping. Two or more individuals possess shared understanding when they have exactly the same kind of understanding, relative to a particular task context or epistemic domain. We might expect to see this kind of understanding in cases where individuals have exactly the same role with respect to a given task. Note that this view of understanding does not require agents to possess the same understanding across all domains and tasks; it simply requires them to have the same understanding within a particular domain or task.

5.2. Similar Understanding

Unlike complete understanding, the notion of similar understanding does not assume that two or more individuals need to have identical forms of understanding. Instead, what is deemed important is similarity of abilities. Individuals who have high levels of shared understanding will have similar abilities when it comes to some task, but differences in both the scope (e.g. coverage of particular domains) and depth of understanding present an upper limit on the degree of overlap in understanding. This view of shared understanding is likely to be the most suitable for real-world environments. It treats shared understanding as a variable quantity – something that can exist to a greater or lesser extent.

5.3. Complementary Understanding

In some situations it may not be necessary for agents to have identical, or even similar, abilities in order to successfully contribute to collaborative tasks. What might be required instead is an understanding that is complementary or compatible to whatever understanding is possessed by other group members. One might expect to see this form of understanding in teams featuring a high degree of role specialization (where individual team members are required to perform specialized tasks that are not shared with other team members). It is not compulsory to treat complementary understanding as a form of shared understanding; however, it may be regarded as such if we interpret 'shared' to mean the distribution of understanding across multiple group members. If we do this, however, it becomes somewhat difficult to distinguish complementary forms of understanding from distributed forms of understanding (see Section 5.4). As such, the status of complementary or compatible forms of understanding is currently unclear.

5.4. Distributed Understanding

In Section 4 it was suggested that our notions of shared understanding should be guided by functionalist intuitions. We should not, it was suggested, commit ourselves to a view that accords special significance to one particular mechanistic realization of understanding (e.g. a neural one). Rather, we should allow for the possibility that multiple mechanisms may contribute to the expression of behaviours warranting talk of an agent as understanding some target domain of discourse. It was this commitment to a functionalist viewpoint that enabled us to see machine agents as capable of understanding (see Section 4). In this section, we develop these functionalist intuitions further. We suggest that by allowing notions of understanding to 'float free' of the specific details of physical implementation, we can adopt a view of understanding that sees it as (at least potentially) grounded in processing loops that extend beyond the boundaries of individual agents to incorporate elements of the broader social and technological environment.

To make this idea somewhat more concrete, imagine two teams of human subjects each engaged in a task that requires the explanation and prediction of system states. For the sake of argument, let us say that the target system (the one whose behaviour is being explained and predicted by the teams) is the behaviour of a specific tribal group in a military conflict zone. We assume that predictive and explanatory success in this domain depends on the availability of large and heterogeneous bodies of specialist knowledge (perhaps subtending disciplines as diverse as cultural anthropology, psychology, history, sociology and so on), and it is therefore necessary for the two teams to be composed of experts from multiple disciplines. The objective of each team is to use the expertise of team members in order to support the generation of reliable predictions and explanations regarding tribal behaviour. To the extent that the teams generate similar explanations and predictions, we may conclude that they (i.e. the teams)

possess shared understanding of the target domain[10]. But note that within each team the mechanisms that contribute to understanding are essentially distributed; the mechanisms undergirding explanatorily- and predictively-potent (team-level) outcomes are realized by processing loops that extend across all team members. It is in precisely this way that understanding (or at least the mechanisms that contribute to understanding) can be distributed. Specific individuals within a team need not understand a complex system in order for team-level understanding to emerge, and it may even be the case that shared understanding between team members is somewhat limited – each member may possess unique bodies of expertise and knowledge that is not shared with other members of the same team. What is important to note here is that the mechanisms that contribute to understanding need not be localized to the heads of individual human agents; instead, they may extend across a variety of biological and (sometimes) non-biological resources. In addition, the system that does the understanding may not necessarily be an individual human agent; instead, it may sometimes be a collection of (often) heterogeneous elements comprising multiple human agents, intelligent software systems, networked information resources and other technological artefacts.

Analogues to this distributed approach to understanding are apparent in the literature on SMMs. A core claim in the SMMs literature is that greater inter-individual similarity of mental models leads to greater similarity in the expectations and explanations generated by individuals. This, in turn, is seen to result in improved coordination, communication and other team behaviours [4]. Differences in mental models are expected to result in differences in expectations and explanations. Thus, the best way to ensure optimal team performance is to ensure that individuals possess the same mental models. In contrast to this view, Banks and Millward [15] suggest that individuals do not need to share mental models (at least in the sense that individuals possess multiple similar models). Instead they propose that the cognitive process of running a mental model can be divided or distributed amongst team members. One benefit of this approach, they argue, is that it avoids the need for team members to possess redundant bodies of knowledge. By distributing the workload for running the mental model, team members effectively spread the load imposed on the group; they essentially avoid the effort and work required for individuals to assimilate, maintain and execute full mental models of the target system [see 16, 17].

6. ISSUES FOR SHARED UNDERSTANDING RESEARCH WITHIN COALITION CONTEXTS

The majority of this paper has been devoted to the derivation of definitions for the notions of understanding and shared understanding. The current section seeks to raise some specific questions that should be addressed by future research aimed at improving shared understanding in military coalition contexts.

6.1. Who (or what) shares the understanding?

We saw that our notion of shared understanding makes no commitment about the nature of realizing mechanisms (two agents could possess the same or similar understanding without using the same mechanisms). What implication does this have for our notions of who (or what) is deemed to possess shared understanding? Must it always be the case that notions of shared understanding are developed with regard to individual human agents, or can collections of human (and perhaps machine) agents also manifest shared understanding?

Besides the question of whether shared understanding is something possessed by individuals, groups or machine entities, there is also the question of who *should* share understanding in particular task contexts. Military coalitions have complex group structures and group dynamics. Small ad hoc teams may be assembled in response to specific challenges, and large scale operations may necessitate complex patterns of inter-team coordination and collaboration. Such complexity raises questions about how we should identify which subset of individuals needs to possess shared understanding, and what kind of shared understanding (e.g. scope and depth) they should possess.

6.2. What is the value of shared understanding?

A key question for future research in coalition contexts is the relationship between shared understanding and group performance outcomes. Shared understanding may be important for the accuracy, quality, volume and timeliness of task outcomes, and future research should clearly aim to explore this possibility. Another effect of shared understanding may be to enhance team processes or team behaviours. For example, shared understanding may improve inter-agent communication, enabling both human and synthetic agents to anticipate the information requirements of teammates and provide them with advance information. This ability to proactively provide information in advance of its actual use has been identified as a key aspect of team effectiveness [18, 19][11]. Shared understanding may also improve the efficiency of communication or reduce the need for communication altogether. This can be useful in terms of reducing the burden on communication systems that may be limited in terms of their available bandwidth and power. Finally, it is possible that greater levels of shared understanding may deliver a number of beneficial 'psychoaffective' outcomes. These could include things such as improved morale, trust and team satisfaction.

[10] Similarity of response output is sufficient for conclusions about shared understanding in this case because we assume no interdependence between the teams with respect to the accomplishment of particular goals. Although similarity of response output may be sufficient for conclusions about shared understanding in cases such as these we maintain that similarity of response output is not necessary for two or more agents to possess shared understanding (see Section 3 for more on this).

[11] For example, Oser et al [19] found that the teamwork behaviour of 'offering information before it was requested' was related to team effectiveness in military command and control teams.

6.3. Is shared understanding always desirable?

Even if shared understanding can be found to exert a positive effect on performance in some situations, it is by no means clear that we should strive to enhance shared understanding in *every* situation. There are clearly some situations in which shared understanding will be difficult to establish (e.g. multi-disciplinary teams of experts that are brought together to collaboratively resolve a complex problem). Moreover, in some situations shared understanding may stifle creativity or contribute to negative group behaviours such as groupthink [20].

6.4. What factors contribute to shared understanding?

Inasmuch as SMMs can be considered as one of the realizing mechanisms for shared understanding then some of the interventions that have been proposed to foster the development of SMM may be important in the search for ways to enhance shared understanding [18, 21]. It should be remembered, however, that many of these interventions are based on small-team situations and they may not be suitable for the kind of environments in which coalition operations are typically undertaken. In addition, coalition operations feature a diversity of groups differing with respect to factors such as entitativity, permeability, size, culture and opportunities for interpersonal (face-to-face) interaction. This suggests that the factors contributing to shared understanding may be highly heterogeneous, and that interventions aimed at enhancing shared understanding will need to consider the nature of groups, as well as the dynamics of inter-group interaction.

6.5. What kind of understanding is required?

The kind of understanding that needs to be shared by two or more agents will vary depending on the nature of the tasks in which the agents participate, as well as the nature of the agents themselves (e.g. whether they are all humans, software agents or some mixture of the two) [22]. For example, shared understanding may target aspects of a particular situation, the dynamics of team interaction or the strategies that need to be pursued in order to realize task goals. Research thus needs to be sensitive to the kind of understanding that is required in particular situations.

It should also be remembered that multiple forms of shared understanding may be conceptualized (see Section 5), and that not all these forms of shared understanding require agents to posses similar or overlapping abilities. In some task contexts, it may make more sense to adopt a more distributed perspective with respect to shared understanding (see Section 5.4).

6.6. How should shared understanding be measured?

It has not been the purpose of this paper to consider approaches to the measurement of shared understanding; nevertheless, any scientific progress on the issue of shared understanding obviously requires reliable measurement techniques. Given our ability-based definition of understanding and shared understanding, the operationalization and measurement of shared understanding needs to focus on the kind of responses that are made by subjects. In some situations it may be appropriate to regard similarity of response output as indicative of shared understanding, although, as noted in Section 3, similarity of response output is not required for shared understanding.

Since SMM may constitute one means by which shared understanding is realized, the kind of techniques used to measure SMMs may have some validity in measuring shared understanding[12]. The validity of these techniques will ultimately be based on the extent to which the similarity of inter-individual mental models predicts the level of shared understanding between those individuals. Since this still requires some independent measure of shared understanding to be formulated, SMM approaches are only likely to be worthwhile if they are easier or more cost-effective to deploy than alternative approaches to the measurement of shared understanding.

Some of the techniques used to measure team knowledge [24] and situation awareness [8, 25, 26] may also be useful in the further scientific exploration of shared understanding.

7. CONCLUSION

In this paper we have sought to explore a number of issues associated with the notion of shared understanding, specifically in relation to military coalition operations. A primary aim of the paper has been to clarify what is meant by the terms ‘understanding’ and ‘shared understanding’, and we have sought to formulate definitions for each of these terms in ways that (hopefully) assists with future empirical studies. Our proposed definition for understanding (i.e. individual understanding) emphasizes the ability of agents to exploit bodies of causal knowledge. In order to understand, we suggest, agents need to be able to exploit causal knowledge in order to adapt response output in goal-relevant ways (i.e. ways that ensure the efficient realization of cognitive and behavioural goals). The main manifestation of this ability (at least in the human case) is the ability to formulate expectations and explanations regarding the part of the world to which understanding applies (i.e. the domain of understanding). Thus in the case of understanding a complex system (perhaps another human agent), understanding is apparent when we are able to reliably predict future states (e.g. behavioural outcomes) and account for such states in terms of causally-significant forces and factors (e.g. the causal role of mental states with respect to the expression of overt behaviour). Such predictive and explanatory behaviour is an example of how we are able to exploit causal knowledge in order to accomplish specific goals (e.g. reliably predict behaviour).

Having proposed a definition for understanding, we suggested that shared understanding is the ability of

[12] Shared mental models are typically operationalized as the correlation between link-weighted networks using Pathfinder analysis [23].

multiple agents to exploit common bodies of causal knowledge for the purpose of accomplishing common (or shared) goals. Our notion of shared understanding emphasizes the way in which two or more agents are able to use causal knowledge to guide thought and action in common ways. Thus shared understanding will be apparent whenever two or more agents are able to use causal knowledge to generate similar explanations and predictions. In the case of understanding intentional action, two agents may be said to have a shared understanding if they are able to make the same actual (and counterfactual) predictions concerning overt behaviour, and they are additionally able to formulate similar explanations (e.g. in terms of the possession of particular mental states) as to why such predictions respect the 'causal logic' of the domain in question.

Both of the definitions that we have proposed see understanding as a kind of ability – a way to exploit information and knowledge in order to realize cognitive and behavioural goals. In part, this definition is inspired by our experience of the way in which shared understanding is deemed to influence coalition military effectiveness. Thus in talking about coalition operations, shared understanding is often cited as a factor that enables coalition force elements (perhaps from different command structures) to adaptively coordinate their collective behaviours in order to accomplish common mission objectives (Maj. Edward Gentle, personal communication). It seems that some sort of predictive ability (perhaps of the situation, group behaviour, or commander decision-making) must underpin this coordinative ability. Thus it is only when force elements are able to anticipate the behaviour of other elements, and predict the effects of actions on the environment, that they are able to coordinate (or synchronize) their collective actions in support of common goals[13].

Inasmuch as shared understanding contributes to the deployment of efficient modes of inter-agent communication and coordination, it may have particular value in the context of coalition operations. This is because coalition environments are often resource constrained environments in which power overheads and network traffic must be kept to a minimum. If shared understanding improves the efficiency of inter-agent communication (perhaps reducing the need for communication altogether), it may optimize the use of limited network resources. In addition, shared understanding may enable coalition members to anticipate one another's information requirements and thereby optimize the distribution of information within a network environment (i.e. sending information to the right place at the right time). This ability to proactively provide information in advance of its actual use has been shown to improve team effectiveness in empirical studies [18, 19], and it is also a major focus of research efforts in the synthetic agent community [27, 28].

[13] Obviously much may also depend on the extent to which force elements have an *awareness* of what other force elements are currently doing.

In conclusion, shared understanding emerges as a potentially important construct in enhancing coalition effectiveness. A number of problems with past coalition operations have been attributed to breakdowns in shared understanding (e.g. Operation Anaconda), and shared understanding is typically cited as a desirable feature of coalition operations. Moreover, major defence-related research programmes, such as the International Technology Alliance[14] (ITA), have identified shared understanding as a hard problem for future coalition operations. In light of this, it is important that the scientific community should strive to generate definitions of shared understanding and posit empirically testable hypotheses that can be used to guide future research. This paper constitutes an initial step in the direction of this goal. It has proposed a specific definition for shared understanding and identified a number of challenges for future research. The work presented here will, we hope, serve to stimulate theoretical debate and guide empirical research regarding the nature and value of shared understanding in military coalition environments.

ACKNOWLEDGMENT

This research was sponsored by the US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

[1] L. Wittgenstein, *Philosophical Investigations*. Oxford, UK: Basil Blackwell, 1967.

[2] J. Cannon-Bowers, E. Salas, and S. Converse, "Cognitive psychology and team training: Training shared mental models and complex systems," *Human Factors Society Bulletin*, vol. 33, pp. 1–4, 1990.

[3] J. A. Cannon-Bowers, E. Salas, and S. Converse, "Shared mental models in expert team decision making," in *Individual and Group Decision Making: Current Issues*, J. N. Castellan, Ed. Hillsdale, New Jersey, USA: Erlbaum, 1993.

[4] W. Rouse, J. Cannon-Bowers, and E. Salas, "The role of mental models in team performance in complex systems," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 22, pp. 1296-1308, 1992.

[5] W. B. Rouse and N. M. Morris, "On looking into the black box: Prospects and limits in the search for mental models," *Psychological Bulletin*, vol. 100, pp. 349-363, 1986.

[6] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," *Human factors*, vol. 37, pp. 32-64, 1995.

[7] J. Mathieu, T. Heffner, G. Goodwin, J. Cannon-Bowers, and E. Salas, "Scaling the quality of teammates' mental models: equifinality and normative comparisons," *Journal of Organizational Behavior*, vol. 26, pp. 37-56, 2005.

[8] M. R. Endsley, "Measurement of situation awareness in dynamic systems," *Human Factors*, vol. 37, pp. 65-84, 1995.

[9] A. A. Nofi, "Defining and Measuring Shared Situational Awareness," Center for Naval Analyses, Alexandria, Virginia, USA 2000.

[14] http://www.usukita.org/

[10] P. R. Smart, A. Bahrami, D. Braines, D. McRae-Spencer, J. Yuan, and N. R. Shadbolt, "Semantic Technologies and Enhanced Situation Awareness," in *1st Annual Conference of the International Technology Alliance (ACITA'07)*, Maryland, USA, 2007.

[11] J. Searle, "Minds, brains, and programs," *Behavioral and Brain Sciences*, vol. 3, pp. 417-424, 1980.

[12] M. Price, "Should We Expect To Feel As If We Understand Consciousness?," in *Explaining Consciousness: The Hard Problem*, J. Shear, Ed. Cambridge, Massachusetts, USA: MIT Press, 1997.

[13] E. Rosch, "Is causality circular?," *Journal of Consciousness Studies*, vol. 1, pp. 50-65, 1994.

[14] J. A. Cannon-Bowers and E. Salas, "Reflections on shared cognition," *Journal of Organizational Behavior*, vol. 22, pp. 195-202, 2001.

[15] A. P. Banks and L. J. Millward, "Running shared mental models as a distributed cognitive process," *British Journal of Psychology*, vol. 91, pp. 513-531, 2000.

[16] K. Sycara and C. M. Lewis, "Modeling Group Decision Making and Negotiation in Concurrent Product Design," *Systems Automation: Research and Applications*, vol. 1, 1991.

[17] C. M. Lewis and K. Sycara, "Informed Decision Making in Multi-Specialist Cooperation," *Group Decision and Negotiation*, vol. 2, pp. 279-300, 2004.

[18] R. Stout, J. Cannon-Bowers, E. Salas, and D. Milanovich, "Planning, Shared Mental Models, and Coordinated Performance: An Empirical Link Is Established," *Human Factors*, vol. 41, pp. 61-71, 1999.

[19] R. Oser, C. Prince, and B. B. Morgan, "Differences in aircrew communication content as a function of flight requirement: Implications for operational aircrew training," in *34th Annual Meeting of the Human Factors Society*, Orlando, Florida, USA, 1990.

[20] I. L. Janis, *Victims of Groupthink*, 2nd ed. Boston, USA: Houghton Mifflin, 1982.

[21] K. A. Smith-Jentsch, J. A. Cannon-Bowers, S. I. Tannenbaum, and E. Salas, "Guided team self-correction: Impacts on team mental models, processes and effectiveness," *Small Group Research*, vol. 39, pp. 303-327, 2008.

[22] K. Sycara and C. M. Lewis, "Integrating Intelligent Agents into Human Teams," in *Team Cognition: Understanding the Factors That Drive Process and Performance*, E. Salas and S. Fiore, Eds. Washington D.C., USA: American Psychological Association, 2004, pp. 203-233.

[23] R. Schvaneveldt, F. Durso, T. Goldsmith, T. Breen, N. Cooke, R. Tucker, and J. Maio, "Measuring the Structure of Expertise," *International Journal of Man-Machine Studies*, vol. 23, pp. 699-728, 1985.

[24] N. Cooke, E. Salas, J. Cannon-Bowers, and R. Stout, "Measuring team knowledge," *Human Factors*, vol. 42, pp. 151-173, 2000.

[25] R. M. Taylor, "Situational awareness rating technique (SART): The development of a tool for aircrew systems design," in *Situational Awareness in Aerospace Operations (AGARD-CP-478* Neuilly Sur Seine, France: NATO-AGARD, 1990, pp. 3/1-3/17.

[26] F. T. Durso, C. A. Hackworth, T. R. Truitt, J. Crutchfield, D. Nikolic, and C. A. Manning, "Situation awareness as a predictor of performance for en route air traffic controllers," *Air Traffic Control Quarterly*, vol. 6, pp. 1-20, 1998.

[27] J. Yen, X. Fan, S. Sun, T. Hanratty, and J. Dumer, "Agents with shared mental models for enhancing team decision makings," *Decision Support Systems*, vol. 41, pp. 634-653, 2006.

[28] J. Yen, J. Yin, T. R. Loerger, M. Miller, D. Xu, and R. A. Volz, "CAST: Collaborative Agents for Simulating Teamwork," in *17th International Joint Conference on Artificial Intelligence (IJCAI)*, Seattle, Washington, USA, 2001, pp. 1135–1142.