Artificial Intelligence and Social Computing, Vol. 28, 2022, 194–205
https://doi.org/10.54941/ahfe1001463

Improving Common Ground in Human-Machine Teaming: Dimensions, Gaps, and Priorities

Robert E. Wray[1], James Kirk[1], and Jeremiah T. Folsom-Kovarik[2]

[1]Center for Integrated Cognition @ IQMRI Ann Arbor, MI 48106, USA
[2]Soar Technology, Inc. Ann Arbor, MI 48105, USA

ABSTRACT

"Common ground" refers to knowledge, facts, beliefs, etc. shared between participants in a joint activity. We seek to enable agents to partner in building and maintaining common ground with human teammates which is an acknowledged need in human-machine systems. This paper provides a detailed analysis of the components of common ground. We present a simple taxonomy of the individual components and dimensions of common ground, map components to specific classes of functions (agent capabilities required to build and maintain common ground) and identify how deficits will manifest (types of errors that may arise when the functions are insufficient for a particular component). We explore how mentalizing, a cognitive capability that allows an entity to reason about the beliefs, goals, and intents of others, may be useful for mitigating a number of the acute issues identified in the analysis.

Keywords: Human-machine Teaming, Common ground, Perspective-taking

INTRODUCTION

Common ground refers to the knowledge, facts, beliefs, etc. shared between participants in some joint activity (Clark, 1996). We seek to enable human-machine teaming agents to more fully partner in building and maintaining common ground with their human teammates. Enabling agent capability includes computational grounding (Taylor, 2017; Traum, 1994), initiating common ground repairs and responding to requests for repair, and adapting one's understanding of a joint activity.

Common ground is an identified gap for human-machine teaming (Mouloua and Hancock, 2019). This paper provides a detailed analysis of the components of common ground relevant to this gap. We present a simple taxonomy of the individual components and dimensions of common ground; map components to specific classes of agent functions or capabilities required to build and maintain common ground; and identify how deficits will manifest – types of errors that may arise when the functions are insufficient for a particular component of the common ground.

We scope the analysis to human-machine teaming in a military context, where interactions are task-oriented and human participants are typically well-trained. We use teammates to emphasize that the participants in the joint

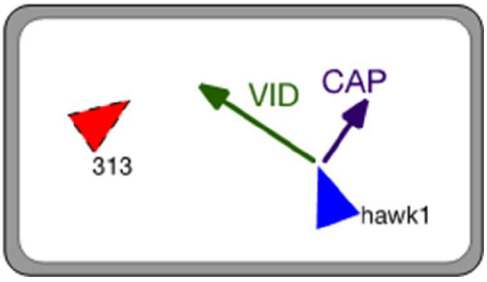

Figure 1: Human-machine teaming in an air defense scenario.

activity share training and a common purpose. We also scope to the meaning-and-understanding level of communication. Lower levels of communication, such as forming utterances and vocalizing them, are important in building common ground in real-world interactions (i.e., one may misspeak or mishear some vocalization). However, focusing on meaning results in an emphasis on potential misalignments in either teammate's conceptual understanding, which is our primary emphasis.

The analysis thus defines how specific deficits in shared knowledge or processing differences manifest in misalignment in shared understanding. The paper identifies specific challenges and prioritizes them according to acuteness of need. Not all of the gaps require immediate attention to improve human-machine interaction (HMI). Further, the solution to specific issues may sometimes depend on solutions to other issues. As a consequence, this analysis facilitates greater understanding of how to attack issues in misalignment in both the nearer- and longer-terms.

MOTIVATING EXAMPLE

In order to illustrate the components of common ground, we introduce a simple example, illustrated in Figure 1. This scenario is inspired by the influential Tactical Decision Making Under Stress (TADMUS) effort (Cannon-Bowers and Salas, 1998). TADMUS researched how human teams coordinate their behavior in high stress, uncertain tactical situations. Our review of common ground issues evident in human-human interactions in TADMUS informs this analysis.

In this scenario, a human operator and an agent are working jointly to monitor threats in an airspace. The agent has a pre-defined set of activities it can perform autonomously, such as assigning a friendly (blue) asset to perform a visual identification (VID) task. The operator can take on any task the agent can perform, can override agent decisions, and can assign some tasks that the agent does not have authority to perform (e.g., a lethal intercept).

In the figure, the agent has concluded that the appropriate task for the "hawk1" entity is to perform a VID of track 313, a potentially hostile aircraft. However, the operator has concluded that hawk1 should instead move to a different task, a combat air patrol (CAP). This disagreement may serve to

surface differences in the underlying common ground. Why might such differences arise, and how do these differences relate to building and maintaining common ground? In the next section, we use this example to explore answers to these questions.

COMPONENTS OF COMMON GROUND

We build on Clark's (1996) framework. He describes three temporal divisions and personal or communal bases for establishing common ground. We relate these divisions to the cooperative HMI task in Figure 1, to identify which have the greatest impact in the context of human-machine teaming in military applications.

Temporal Divisions of Common Ground

Does a basis for believing something is in the common ground come from some experience or learning in the collective pasts of the teammates, from recent interactions, or from the present moment in which they share a common or similar perception of the situation? These are the temporal divisions of common ground.

Initial Common Ground: Initial common ground derives from assumptions, beliefs, knowledge, procedures, etc. that each participant can (generally) assume are known to the other. In a military context, where thorough training is typically emphasized, significant commonality in prior knowledge among human participants in a joint activity can be assumed; in such a situation, the initial common ground maps to shared mental models (Cannon-Bowers et al., 1993). In Fig 1, we can assume that both the agent and human understand the different tasks that could be assigned to hawk1, like VID and CAP. Although shared training provides a solid foundation for task-focused joint action, many issues will still arise due to differences in capacities, background and experience even with common training and experiences.

Public Events so Far: The public events so far are a record of what has happened in recent interactions between entities. If the operator told the agent that track 313 was not hostile, then both the fact (the track is not hostile) and the source of the assertion (operator designated the track as not hostile) should be added to the common ground. For HMI, issues may arise when there is not agreement on which recent interactions are public or when a list of events becomes burdensome for the entities to maintain.

Current State of the Activity: The current state encapsulates what the participants apprehend as the current situation. It includes direct perception but also simple entailments of it. For example, if track 313 was headed directly toward hawk1 and moving very fast, that might entail a higher threat estimation. Misalignment can readily arise when teammates have different vantages on a situation. Further, by their very natures, humans and machines bring different perspectives on seeing and interpreting the current state, a frequent source of common-ground issues in HMI.

Personal Common Ground

Personal common ground concerns the shared or joint experiences that entities have with one another. The components of personal common ground include perception, action, and history established in interactions to date. Personal common ground relates especially to the current moment (current state of active and public events so far) but can include past history as well. We outline how issues in understanding or interaction may arise from misalignment in common ground in those dimensions.

Perceptual Bases

Perceptual bases are what is shared across the perception of entities in an interaction.

Initial Common Ground: For trained military operators, deviations in initial common ground attributable to perception seem negligible. One purpose of training is to ensure that teammates have comparable mental models to see situations similarly. However, in an HMI context, differences in the core structure of perception could lead to misalignment as well. For example, if a human teammate makes a rule to handle "fast-moving threats" and the agent lacks a definition of this somewhat amorphous category, then the difference in perceptual basis might be problematic.

Current State of Activity/Public Events so Far: Differences in perspective can obviously lead to differences in understanding of the current state. The literature on human communication offers many examples of such misalignment (Clark and Brennan, 1991; Pickering and Garrod, 2004). Does a reference to "the threats on the left" describe the left side of the screen or the left of the friendly interceptor? Differences in vantage can also lead to differences in an understanding of the public events so far. For example, if an agent and operator settle on an object as a common reference point and later discover that they did not choose the same object as a reference point, then the public events so far would need revision.

Actional Bases

Actional bases capture the understanding of what joint actions are available and have been undertaken. With training, common ground issues arising from actional bases should be relatively rare (unless they derive from issues elsewhere). For example, misaligned perception of the situation or understanding of procedures and techniques could result in misunderstandings about which joint actions are applicable or relevant.

Personal Diaries

A personal diary is an autobiographical memory (Tulving, 1983) of what has occurred between teammates over time. A personal history is a record of the public events so far that occurred between some set of actors continuing back over many interactions. Expectations and patterns that arise from repeated interaction are a clear component of teamwork.

Initial Common Ground: In human-human interaction, prior shared experience (personal diaries) can be assumed to be part of the common ground

between those actors. This is not true in today's HMI and thus limits the interaction. The agent cannot interpret a statement akin to "Do you recall when we had an issue like this yesterday and we resolved it by ...." This lack of awareness of past interaction could lead to discrepancies in the common ground (the human can refer to it; the agent cannot) but also frustration of the operator (the agent cannot recall, build on, or improve from the prior interaction).

Current State of Activity/Public Events so Far: A memory of previously experienced states/events so far and mapping to the current situation would generally be useful for an agent to shortcut extra reasoning about meaning and references.

Communal Common Ground

Communal common ground refers to bases for common ground that do not arise from direct prior experience or interaction between conversants. Two participants in an interaction may come from various cultural contexts and those contexts often need to be uncovered in the conversation. Cultural context is however highly constrained by teamwork in a military setting. Prior training sets an expectation for familiarity with the culture of the branch, warfare area, and unit between participants conducting a joint activity such as coordinated air defense in the Figure 1 example.

We enhance Clark's accounting of communal common ground in three ways:

1. We circumscribe Clark's notion of *human nature* to common sense, or background knowledge that most human adults will know. This distinguishes human nature from task-specific facts, procedures, and norms.
2. Clark groups together facts, procedures, and norms as one category in the communal common ground. Because implications for an agent are somewhat different for each of these, we break them out into distinct categories.
3. The context of military, task-oriented conversation lets us consider knowledge of facts, procedures, and norms arising from the task domain. Issues arising from differences in communal common ground would be broader and more demanding for agents when a shared task context cannot be assumed.

Human Nature

Clark calls out human nature as commonsense understanding such as folk psychology (Bratman, 1987) and naïve physics (Hayes, 1979). In human interaction, conversants share some understanding of how the world and people in it generally act, which then informs building and maintaining common ground.

We extend this notion of human nature to include some understanding of human processing limitations and biases (Tversky and Kahneman, 1981). While individuals may not have explicit knowledge of these phenomena, they generally understand, one cannot generate a long list of items and expect the

hearer to internalize it after a single presentation, that some detail mentioned a few minutes ago may not remain in the other's mind, and that people generally have other predictable limitations. Failure to model these limitations in agents used for HMI may lead to issues in building and maintaining common ground with a human who manifests such limitations.

Initial Common Ground: A lack of common-sense reasoning available to an agent can lead both to tedious interactions (human explains every detail) and to gaps in common ground (when an assertion is not sufficiently elaborated for either agent or teammate to apprehend that that there is misalignment). For example, suppose the human has assigned the CAP in Figure 1 because the human is aware that track 313 has been targeted and assumes it will be neutralized. The agent might misunderstand the human to mean that the entity is not actually hostile if it only sees the CAP order.

Current State of Activity/Public Events so Far: Differences in attention and attention management are frequent sources of difference in understanding or apprehending the current state. Issues arise both because attention can be elsewhere and because the state can be complex, with too much information for a human operator to take in. Agents often have more capability to comprehend a complex state than humans do. Humans generally have a greater ability to attend to what is important or salient in a complex environment and also to create abstract representations that elide some detail and complexity. In Figure 1, the operator might assign the CAP to hawk1 because they are unaware of track 313.

Domain Facts

Basic factual knowledge should be readily shared in an environment with high levels of training. For HMI, common ground issues arising from facts alone are likely to be rare unless some relevant domain facts are not accessible to the agent.

Domain Procedures

The procedural knowledge within a military task domain includes TTPs (tasks, techniques, and procedures) and application of doctrine. Training generally enables individual and joint understanding of these procedures. For HMI, issues may arise if agent and human understanding of procedures are mismatched.

Initial Common Ground: Whenever participants do not share the same understanding about how to perform their joint activity, misunderstandings are likely to arise. If new tactics say to ignore potential threats that are not on a clear threat vector (pointed toward blue assets), but the agent's knowledge still reflects earlier doctrine to always carry out a VID when an intercept is at all possible, a conflict in initial common ground arises.

Current State of Activity/Public Events so Far: Differences in understanding domain procedures can result in differences in the current understanding of the state. For example, the human and agent may reach different conclusions about the level of threat of track 313, resulting in different orders for hawk1.

Domain Norms

Norms influence human behavior. Norms arise and vary across cultural contexts (Gibbs, 1965). Norms are difficult for agents because, unlike factual and procedural knowledge, norms typically lack explicit codification (Conte and Castelfranchi, 1999). That is, because procedures and facts are "written down," there is greater possibility of agreement and alignment with them coming into a new interaction. Norms present more difficultly to align because they are (largely) implicit.

Cultural Communities

Domain norms, procedures, and facts are contingent on a shared cultural context. More basic lexical issues may manifest as well (*truck* vs. *lorry*; *in* vs. *on* a street). Common training will generally mitigate cultural differences. However, in military contexts such as joint and multinational exercises and operations, training is not fully shared, and cultural differences will likely emerge. Agents need the ability to rapidly adapt and accommodate to new cultural contexts.

NEED ASSESSMENT

Figure 2 summarizes the characteristics of common ground. The vertical dimension reflects the temporal divisions of common ground. The horizontal dimension breaks out the components of personal and communal ground. Specific misalignments in common ground will arise due to discrepancies in representation, knowledge, and/or processing mappable to these elements. For HMI in military operations, challenges may arise across any of these cells, as the previous section detailed.

The specific conditions under which an interaction occurs can be used to anticipate certain kinds of problems. Imagine an inexperienced traveler, unfamiliar with the local language, trying to arrange for transportation to their hotel. The likely lack of common understanding of procedures and norms and minimal overlap in the cultural communities of traveler and transport operator will provide a weak basis for an initial common ground. Analogously, a common design approach in common-ground human experiments is to separate participants and have them perform novel tasks (jointly creating tangrams, for example). This design limits participants' ability to build from an initial common ground in the task and forces different perceptual bases.

Shifting focus to HMI, different intersections within these dimensions are likely to occur more commonly and/or cause more issues than others. An agent designed to support air defense operations is not likely to face significant issues arising from differences in cultural communities because operator training and agent design create a common cultural basis within the context of the mission. Similarly, while it might be preferable for an agent to develop personalized lexicons for interaction with specific individuals, given the nature of a prescribed task and controlled languages used within them, the benefit of a personal lexicon would be negligible.

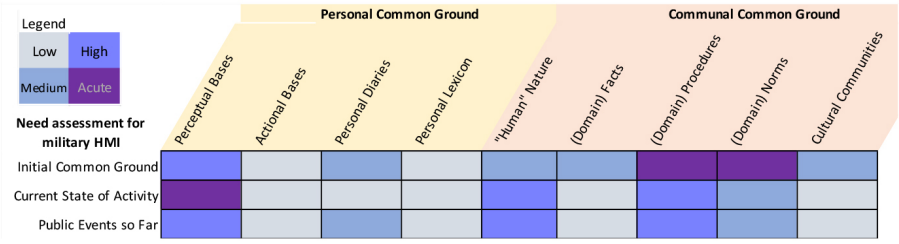

Figure 2: Dimensions of common ground and acuteness of need.

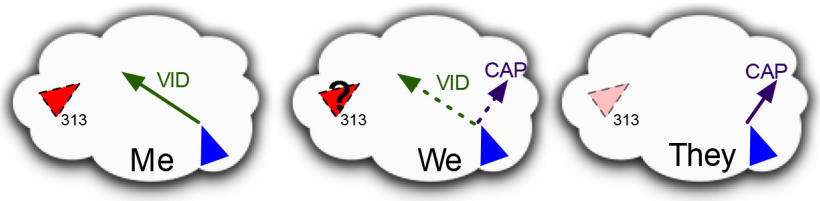

Figure 3: The mentalizing model represents agent understanding of the common ground.

In the figure, we rank (by color) the severity of underlying issues at each intersection among the dimensions. The rankings estimate how critical the need is at that intersection of the dimensions for improved HMI capability to build and to maintain common ground. Rankings consider both capability gaps in today's HMI systems as well as an estimate of the payoffs in addressing those gaps. The rankings draw on our decades of experience in creating HMI systems for the military, consultation with other HMI experts, and formal documents such as procurement requirements.

Based on the rankings, we identify two primary axes of challenge. First, a *weak basis for initial common ground* is frequently problematic across both personal and communal common ground. Training reduces some frictions related to building common ground. However, any military environment will change frequently, making it more of an acute issue to ensure tactics (procedures) and norms are shared.

Second, *core differences in perceptual bases* between human and machine are problematic for common ground. Thus, they will often need to negotiate through interaction a mutual understanding of the current state. The lack of a common perceptual basis is compounded in HMI because the human and agent interact with the shared environment using distinct and largely foreign vantages not only in terms of communication, but in every aspect of perception, from sensory faculties to memory, attention, and even different mental processing.

THE POTENTIAL ROLE OF MENTALIZING

How might these challenges be addressed? Human beings cannot know what is in someone else's head; meaning cannot be directly conveyed. However,

adult humans (and some primates) have a (seeming) innate capacity to imagine what that other person sees and understands (Frith and Frith, 2006). Further, a person can attempt to "put themselves in the shoes of" another person, using their observations of the behavior, communication, etc. of the other to attempt to comprehend (at some level of abstraction and imperfection) the mental states of the other (Nickerson, 1996). This capability is known variously as mentalizing, theory of mind, and perspective-taking. For simplicity, we use *mentalizing* to refer to this capability. A recent review posits that mentalizing is necessary for improved HMI (Kopp and Krämer, 2021).

We hypothesize that both of the challenges identified above can be mitigated by providing agents with greater ability to model and to reason about teammates as outlined by Kopp and Krämer (2021). We seek to enable the agent to model what it knows, believes, assumes (*Me*), what another teammate likely knows, believes, assumes (*They*) and based on these models, to construct a model of what the agent believes is common ground between itself and the other (*We*).

Figure 3 sketches what the proposed extension to agent capabilities may provide an agent. In addition to its pre-existing internal representations of the task, the agent will explicitly represent its conception of the common ground between it and a teammate. (For simplicity, we describe a single teammate, but the overall approach can be applied to a joint activity with more than two actors.) This conceptualization of the common ground includes three components:

Me: The agent creates a simplified representation of its own understanding. This representation includes the current state and public events so far – prior discourse and actions in the joint activity. The *Me* representation includes pointers to task knowledge the agent used to draw conclusions or assertions. The initial common ground is not explicit in *Me*. However, assertions that derive from facts, procedures, and norms the agent uses to reason about the task will refer to these initial common ground dependencies so the agent can reason about them.

For the example, the agent observes track 313. It creates an intention (green arrow) for hawk1 to VID the red aircraft.

They: The agent also builds a simplified representation of what it hypothesizes its partner in the joint activity (e.g., the operator in the air defense mission) knows. This simplified representation employs the agent's task knowledge but applies it to the estimated perspective of the operator. A key requirement, as noted above, is that the agent has an ability to model the operator's cognitive processes and human differences from a machine agent such as attention and memory decay.

In the figure, the operator indicates an intent to assign hawk1 to a CAP. The agent's modeling of the operator enables it to recognize uncertainty about whether the operator has attended to the red aircraft. The task assignment introduces uncertainty, even if the aircraft is visible on the operator's user interface. The pink color of track 313 is intended to convey this uncertainty in the figure. From the agent's perspective, the intention of the operator to assign hawk1 to a more routine and less urgent CAP task makes it more likely that the operator may not be aware of the red aircraft.

We: The agent builds and maintains a joint representation that captures agreement and potential disagreement between *Me* and *They* representations. A key role of this representation is to enable the agent to detect and prioritize discrepancies. In the figure, there is misalignment about the task assignment, but the agent can also readily recognize the potential discrepancy in the awareness of the red aircraft as well.

CONCLUSION

This paper summarizes a fine-grained examination of the components of common ground, focusing on potential issues that may arise in meaning or understanding in human-machine interaction in a military context. The analysis identifies how specific classes of issues may manifest due to underlying misalignments in these various components of the common ground. We prioritized these issues according to the acuteness of need in future HMI systems.

We identified two dimensions of acute need: a more robust initial common ground between agent and human and greater alignment in perceptual bases. We hypothesize that imbuing agents with the ability to reason about the mental states of teammates offers potential to address the acute needs we identified.

Weak basis for initial common ground: Reasoning about the understanding of teammates and reconciling that understanding with the agent's own model of a situation will help surface discrepancies in initial common ground faster and more efficiently. Faster resolution is important because earlier repair reduces downstream errors. Errors can have catastrophic outcomes if they remain unaddressed. Efficiently acting to repair the common ground only as needed is important because time pressure is typical in the performance of most operational military tasks.

Conversational repairs can be used to work around misalignments in the moment. However, without changes to the underlying knowledge that led to a discrepancy in the initial common ground, the problem will recur. We envision mentalizing used in combination with *interactive task learning* (ITL, Gluck and Laird, 2019). Mentalizing contributes to effective repair in the moment. However, when the human and agent are executing a mission, time pressure and the cadence of the activity may prevent immediate reconciliation of underlying differences. After the mission is completed, ITL provides the ability for an operator to explain how the team can improve long-term alignment of their knowledge to avoid similar issues in the future. Mentalizing aids ITL by identifying salience and location of a discrepancy; these are cues for a subsequent interaction to resolve differences.

Core differences in perceptual bases between man and machine: Explicitly modeling a human teammate's mental processing will enable an agent to act proactively and economically to reduce discrepancies in the common ground. Has the human in Figure 3 missed the presence of track 313, its location, or another fact that influences the action? The current situation, what has happened so far, and the agent's underlying knowledge can be used to evaluate

alternatives and pinpoint a source of misalignment between its understanding and the operator's (apparent) understanding of the current situation.

Proactive mitigation should lead to fewer overall repairs, which will reduce operational errors that can arise from misalignments. Economy is important because there will likely be many potential discrepancies that an agent might address – so many that repairs of every misalignment could overwhelm an operator. Instead, mentalizing allows the agent to be deliberate about which and when apparent misalignments in the common ground need to be repaired.

ACKNOWLEDGMENT

This material is based upon work supported by the United States Air Force under Contract No. FA8650-22-P-6412. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force.

REFERENCES

Bratman, M.E., 1987. Intentions, Plans, and Practical Reason. Harvard University Press, Cambridge, MA.

Cannon-Bowers, J.A., Salas, E., 1998. Decision making under stress: Implications for training and simulation. American Psychological Association, Washington, DC.

Cannon-Bowers, J.A., Salas, E., Converse, S., 1993. Shared mental models in expert team decision making, in: Individual and Group Decision Making: Current Issues. Lawrence Erlbaum Associates, Inc, Hillsdale, NJ, US, pp. 221–246.

Clark, H.H., 1996. Using Language. Cambridge University Press.

Clark, H.H., Brennan, S.E., 1991. Grounding in Communication, in: Resnick, L.B., Levine, J.M., Teasley, J.S.D. (Eds.), Perspectives on Socially Shared Cognition. American Psychological Association.

Conte, R., Castelfranchi, C., 1999. From conventions to prescriptions. Towards an integrated view of norms. Artificial Intelligence and Law 7, 323–340. https://doi.org/10.1023/A:1008310107755

Frith, C.D., Frith, U., 2006. The Neural Basis of Mentalizing. Neuron 50, 531–534. https://doi.org/10.1016/j.neuron.2006.05.001

Gibbs, J., 1965. Norms: The problem of definition and classification. American Journal of Sociology 70, 586–594.

Gluck, K., Laird, J. (Eds.), 2019. Interactive Task Learning: Agents, Robots, and Humans Acquiring New Tasks through Natural Interactions, Strüngmann Forum Reports. MIT Press, Cambridge, MA.

Hayes, P., 1979. The naive physics manifesto, in: Expert Systems in the Micro-Electronic Age. Edinburgh University Press, Edinburgh.

Kopp, S., Krämer, N., 2021. Revisiting Human-Agent Communication: The Importance of Joint Co-construction and Understanding Mental States. Front. Psychol. 12, 580955. https://doi.org/10.3389/fpsyg.2021.580955

Mouloua, M., Hancock, P. (Eds.), 2019. Human Performance in Automated and Autonomous Systems: Current Theory and Methods. CRC Press.

Nickerson, R.S., 1996. How we know—and sometimes misjudge—what others know: Imputing one's own knowledge to others. Psychological Bulletin 125, 737–757.

Pickering, M.J., Garrod, S., 2004. Toward a mechanistic psychology of dialogue. Behav. Brain Sci. 27. https://doi.org/10.1017/S0140525X04000056

Taylor, G., 2017. Natural, Multi-modal Interfaces for Unmanned Systems, in: Engineering Psychology and Cognitive Ergonomics: Cognition and Design: 14th International Conference, EPCE 2017. Springer International Publishing, Germany.

Traum, D.R., 1994. A computational theory of grounding in natural language conversation. (Dissertation). Rochester Univ NY Dept of Computer Science.

Tulving, E., 1983. Elements of Episodic Memory. Oxford University Press.

Tversky, A., Kahneman, D.J., 1981. Judgment under Uncertainty: Heuristics and Biases. Science 185, 1124–1131.