



AFRL-RH-WP-TR-2023-0032

ACTIVELY BUILDING CONTEXT (ABC): A MODEL FOR CONTEXT-AWARE COMMUNICATION

**Kole Norberg, Kevin Gluck, Karla Seale
Resilient Cognitive Solutions LLC
113 South 26th Street, Suite 200
Pittsburgh, PA 15203**

January 2023

Interim Report

DISTRIBUTION STATEMENT A: Approved for public release; distribution unlimited.

**AIR FORCE RESEARCH LABORATORY
711 HUMAN PERFORMANCE WING
AIRMAN SYSTEMS DIRECTORATE
WARFIGHTER INTERACTIONS & READINESS DIVISION
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

Qualified requestors may obtain copies of this report from the Defense Technical Information Center (DTIC).

AFRL-RH-WP-TR-2023-0032 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

GREGORY FUNKE, DR-III, Ph.D.
Work Unit Manager
Collaborative Technologies Branch
Airman Systems Directorate
711th Human Performance Wing
Air Force Research Laboratory

TERRY STANARD, DR-III, Ph.D.
Chief, Collaborative Technologies Branch
Airman Systems Directorate
711th Human Performance Wing
Air Force Research Laboratory

LOUISE A. CARTER, DR-IV, Ph.D.
Chief, Warfighter Interactions and Readiness Division
Airman Systems Directorate
711th Human Performance Wing
Air Force Research Laboratory

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE January 2023	2. REPORT TYPE Interim	3. DATES COVERED	
		START DATE 01 November 2021	END DATE 01 December 2022
3. TITLE AND SUBTITLE Actively Building Context (ABC): A Model for Context-Aware Communication			
5a. CONTRACT NUMBER FA8650-20-D-6207	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER	
5d. PROJECT NUMBER	5e. TASK NUMBER	5f. WORK UNIT NUMBER H13S	
6. AUTHOR(S) Kole Norberg, Kevin Gluck, Karla Seale			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Resilient Cognitive Solutions LLC 113 South 26th Street, Suite 200 Pittsburgh, PA 15203			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory 711th Human Performance Wing Airman Systems Directorate Warfighter Interactions and Readiness Division Wright-Patterson Air Force Base, OH 45433		10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RH-WP-TR-2023-0032
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION A. Approved for public release; distribution unlimited.			
13. SUPPLEMENTARY NOTES AFRL-2023-3397; Cleared 17 Jul 2023			
14. ABSTRACT Effective communication is important in nearly all aspects of our lives, and it is especially crucial for military missions. Unfortunately, communication failures are common. The effects of these range from confusion, frustration, and delay to consequential losses of assets and lives. Context misalignment is a common factor in communication failures. It may be possible to avoid those failures and their negative impacts, and indeed to improve mission performance, through focused research on context-aware communications. Thus, the goal of the current line of research is to enhance operational communication by improving contextual overlap among both human and machine communicators in military environments. We seek to do this by developing a context-aware machine agent capable of dynamically determining the relative importance of information as it pertains to the current situation (i.e., identifying context) in order to accomplish one or more of the following: (a) use context to inform its own comprehension and language production processes, (b) aid humans in identifying and tracking context which may be relevant to a communication, and (c) aid humans in detecting and resolving risks for miscommunication.			
15. SUBJECT TERMS Context, Communication, Natural Language, Human-Machine Teaming			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR
18. NUMBER OF PAGES 41			
19a. NAME OF RESPONSIBLE PERSON Gregory Funke, Ph.D.			19b. PHONE NUMBER (Include area code)

TABLE OF CONTENTS

List of Figures.....	ii
List of Tables.....	ii
1.0 EXECUTIVE SUMMARY	1
1.1 Context-Aware Support for Enhanced Operational Communication	1
1.2 Defining Context.....	1
1.3 The Actively Building Context (ABC) Model of Context Aware Communication	1
1.4 Building a Context-Aware Communications Agent	2
1.5 Conclusion	3
2.0 CONTEXT-AWARE SUPPORT FOR ENHANCED OPERATIONAL COMMUNICATIONS	4
2.1 Defining Context, the Situation, and Awareness	5
2.1.1 Error in the Representation	8
2.1.2 Fluctuations According to Attention and Focus.....	8
2.1.3 Theory of Mind	10
2.1.4 Situational versus Contextual Awareness	10
2.2 Defining Context-Aware <i>Communication</i>	10
2.2.1 Situation Model	11
2.2.1.1 Environmental Context	12
2.2.1.2 Memory Context.....	12
2.2.1.3 Textual Context	12
2.2.1.4 Language Familiarity Context	12
2.2.2 Grounding	13
2.2.3 Miscommunication and Clarifications	14
2.2.4 Team Communication	15
3.0 THE ABC MODEL OF CONTEXT-AWARE COMMUNICATION	17
4.0 BUILDING A CONTEXT-AWARE COMMUNICATIONS AGENT	20
4.1 Context Sensing as a Prerequisite	20
4.2 Context Augmenting Technologies.....	21
4.3 Context Adapting Technologies	25
4.4 Context Resource Exploiting Technologies.....	25

4.5	Related work at Air Force Research Laboratory (AFRL) - The Synthetic Teammate Project	26
5.0	CONCLUSION	27
6.0	REFERENCES	28
7.0	LIST OF ACRONYMS, ABBREVIATIONS AND SYMBOLS	35

LIST OF FIGURES

Figure 1. Zimmermann et al. (2007) Model of Context.	7
Figure 2. ABC Model of Context-Aware Communication.	18

LIST OF TABLES

Table 1. Communication Failures Related to Context-Awareness.	4
--	---

1.0 EXECUTIVE SUMMARY

1.1 Context-Aware Support for Enhanced Operational Communication

Effective communication is important in nearly all aspects of our lives, and it is especially crucial for military missions. Unfortunately, communication failures are common. The effects of these range from confusion, frustration, and delay to consequential losses of assets and lives. Context misalignment is a common factor in communication failures. It may be possible to avoid those failures and their negative impacts, and indeed to improve mission performance, through focused research on context-aware communications. Thus, the goal of the current line of research is to enhance operational communication by improving contextual overlap among both human and machine communicators in military environments. We seek to do this by developing a context-aware machine agent capable of dynamically determining the relative importance of information as it pertains to the current situation (i.e., identifying context) in order to accomplish one or more of the following: (a) use context to inform its own comprehension and language production processes, (b) aid humans in identifying and tracking context which may be relevant to a communication, and (c) aid humans in detecting and resolving risks for miscommunication.

1.2 Defining Context

The most broadly accepted definition of context states, “Context is any information that can be used to characterize the situation of an entity” where “an entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves” and where the situation is “a description of the states of relevant entities” (Dey, 2001; Dey & Abowd, 1999). Context also contains gradations of momentary relevance as information moves in and out of attention (Zimmermann et al., 2007). Information that is currently in mind (*in use*) highlights a special level of importance because it is occupying cognitive resources (Dourish, 2004; Winograd, 2001). From there, information will decline in importance according to its possible usefulness until it is no longer considered context.

1.3 The Actively Building Context (ABC) Model of Context Aware Communication

The need for a new conceptual model of context awareness extends from a lack of clarity surrounding how to identify context across multiple situations and how to assess what contextual information is *shared*. Recognizing what information is and is not shared is a vital part of communication (e.g., Clark, 1996; Clark & Brennan, 1991; Clark & Marshall, 1981; Mustajoki, 2012). It requires an agent to have some theory of mind and represent what unique information they have, what unique information their conversation partner might have, and what information both agents likely have. Humans may use a co-presence heuristic to determine if information is available to both parties (Clark & Marshall, 1981). For example, if both parties were present at the same briefing, they will likely consider that information shared even if they do not know for certain that the other person was listening.

Humans also use a process known as grounding (Clark, 1996; Clark & Brennan, 1991) to ensure information required for comprehension is co-present. When conversants believe their partners may have trouble accessing context or that they may be missing context, they will add additional information to *contextualize* their communication. When communicators fail to do this work, conversation partners can be left confused or may infer their own meanings, opening the door for miscommunication (Mustajoki, 2012; Savitsky et al., 2010).

In our model, we create framework for identifying context, determining the need to contextualize a communication, and using context during comprehension. The model represents what each conversant knows individually, what they know together, and what they *believe* their conversation partner knows. As information is introduced into the conversation, it becomes context *in use*. As it falls out of use, it moves into the space of co-presence. Further, the relevance and the availability of information to all members of a conversation is represented on a gradient of accessibility which we propose is a function of time, frequency of use, and salience.

Briefly, time is the measure of distance from the last moment an entity was mentioned in conversation. At time point zero, all parties in the conversation are actively representing the entity. Spreading activation activates information related to the entity, though what information is active, and its degree of activity will vary from person to person. If no further reference to the entity is made, access to this information will decay. Frequency refers to how often the conversant has activated the information previously. Greater frequency predicts faster retrieval speeds (Anderson, 1993). Salience refers to how attention grabbing the information is and is expected to be an interaction between the baseline perceptual qualities of the information (e.g., large) and individual differences in knowledge.

1.4 Building a Context-Aware Communications Agent

Though increasingly sophisticated, machine agents are not yet able to separate context from the pool of all available information without explicit rules (Schilit et al., 2002; van Kasteren & Vredenburg, 2022). Our model provides a path to using communications to identify context. The representations built from the communications can then be used to augment information available to humans, direct resource exploitation, and adapt the environment (cf. Pascoe, 1997).

Context augmentation applications, like Fitbit, determine who receives what information and how the information is delivered. For example, Fitbit detects user activity and sends step reminders accordingly. However, it is designed for a single function and situation. In our more generalized model, the application is only limited by the types of sensors connected to it. The agent will use the communicated goals as well as information *in use* to inform its gradient of activation and direct its attention to specific sensor data. By using the conversation, the machine agent will know what information in the environment is relevant across multiple situations. Using this functionality and a representation of who has access to what information, it will be able to track discrepancies (e.g., a controller provides an inaccurate time to arrival), understand relational language (e.g., “They’re 100 yards east from my position”), track references across multiple exchanges, determine the level of detail to provide a human based on co-presence calculations, and more.

There are still a few core concerns and difficulties for operational application of this model. First, it has been historically difficult to train natural language processing (NLP) models using radio communications due to low signal to noise ratios and diverse accents. In recent years progress has been made addressing those challenges using end-to-end speech recognition architectures with deep neural networks (Badrinath & Balakrishnan, 2022). A second consideration is that context may be intentionally omitted from operational communications in the interest of security. However, we believe that giving the machine agent access to briefings may help with this issue as context from the briefings can be monitored during the mission.

1.5 Conclusion

The research described here establishes a foundation for a dynamic context-aware agent which utilizes communication to understand the relative importance of information in the world (i.e., to recognize context). The machine agent's representation of context can then be leveraged to facilitate human-to-human and human-machine communication. The core advancement of the model is its generalizability, frame for representing context and perspective, and application of computational models of human memory access to context accessibility during communication. Thus, the current work integrates the wealth of literature on context-aware computing, language processing, and memory architectures to create a new, more adaptable, communications agent for use in military operations.

2.0 CONTEXT-AWARE SUPPORT FOR ENHANCED OPERATIONAL COMMUNICATIONS

The military relies on human-machine teaming (HMT) across multiple domains. As the human-machine relationship expands, there is a growing need for humans to be able to communicate with machines using NLP. The enhanced communication abilities will assist human-machine teams in adapting and coordinating their activity to situations as they arise. Further, all human teams also have periodic and sometimes fatal difficulty in communicating with each other. Thus, a central goal of this manuscript is to highlight how NLP can be improved by enhancing machine context-awareness with the end goal of making human-machine *and* human-human communications more flexible, efficient, and reliable.

Communication is the transfer of information from one agent (human or machine) to another, usually through the medium of language. A consistent theme in the communication literature is that contextual alignment facilitates communication (Clark, 1996; Clark & Brennan, 1991; Fischer, 2012; Kayes et al., 2015; Kopp & Krämer, 2021; Matsumoto & Riek, 2022; Menenti et al., 2012; Paleja, 2022; Shishkov et al., 2018; Winograd, 2001; Zimmermann et al., 2007). Alignment occurs when communicators both have ready access to the same information. When there is a misalignment in the context available to communicators, the risk of delays in communication and miscommunication increase. Thus, the risk of a communication failure in military operations is heightened when communicators have divergent context. This is particularly problematic for human-machine communication because to date, even the most sophisticated autonomous agents struggle to represent and use context during communication (Baddour et al., 2019; Freiman et al., 2018; McNeese et al., 2018).

Divergent contexts can come in many forms, but a common one in the military is communication between deployed tactical personnel and their associated command and control leadership. The divergence stems from teams who are not co-located and who have different expertise, goals, and demands on their cognitive resources. Several communication challenges among such teams have been observed during military training operations (Rothwell et al., 2022). Table 1 provides an overview of the challenges as they relate to context-awareness.

Table 1. Communication Failures Related to Context-Awareness.

Issue	Relationship to Context-Awareness
Transmission of inaccurate information	Context-awareness of the producer was represented inaccurately or was incomplete
Missed messages	Speaker's awareness of the listener's context was inaccurate. Listener may have been on a different channel, not present, not listening, or experiencing low signal-to-noise ratios
Clarifications not requested or made	Context provided by producer was incomplete or inaccurate and one of the comprehenders realized it but did not ask for clarification
Misunderstanding	Context provided by speaker was incomplete, neither party realized it, and the comprehender formed an inaccurate or incomplete understanding.

Note. Issues are summarized from Rothwell et al. (2022)

Although these errors were observed during training, the same types of communication errors played significant roles in real-world military and civilian disasters (e.g., Beck & Cohen, 1994; Estival & Molesworth, 2012; Rakas & Yang, 2007; Skaltsa et al., 2012; Tajima, 2004). One harrowing example in terms of context-awareness was the June 2014 friendly fire incident in which a U.S. Air Force B-1 bomber killed five U.S. soldiers and one Afghan soldier. An internal report of the incident concluded that poor situation awareness and ineffective communication were responsible (Lamothe, 2014). All four of the communication issues in Table 1 occurred. The Joint Terminal Attack Controller (JTAC) gave inaccurate information about troop positions, a nearby mountain range interfered with the radio signal resulting in missed messages, the pilots did not question contradictory information about the enemy position, and there was a misunderstanding concerning what the pilots in the B-1 could see on the ground. We will elaborate on this example throughout the manuscript as we illustrate the principles of context-awareness and how variations of a computational machine agent may have been able to intervene.

By exploring the origins of context-aware communication failures in human-human communication, we gain insight into the importance of improving context-awareness among all communicators (both human and machine). We seek to make these improvements by developing a context-aware machine agent capable of dynamically determining the relative importance of information as it pertains to the current situation (i.e., identifying context) in order to accomplish one or more of the following: (a) use context to inform its own comprehension and language production processes, (b) aid humans in identifying and tracking context which may be relevant to a communication, and (c) aid humans in detecting and resolving risks for miscommunication.

We begin this review by discussing two important concepts within the space of context-aware computing: the context and the situation. As will become clear, *the underlying difficulty in creating a dynamic context-aware machine agent is not in representing information but in using the situation to assess the applicability and importance (i.e., relevance) of that information*. In making this argument, we define the distinction between context and the situation before describing the application of both to communication. We then draft a new model for context-aware communication: the ABC model. Our proposed computational model makes new strides in context-aware communication by providing a path for a computational agent to focus in on relevant knowledge and perceptual information as well as recognizing when the information needs to be shared.

2.1 Defining Context, the Situation, and Awareness

Prior literature has defined context in over 150 different ways (Bazire & Brézillon, 2005; cf. Baddour et al., 2019; Brézillon & Pomerol, 1999; Brown et al., 1997; Chen, 2005; Dourish, 2004; Franklin & Flaschbart, 1998; Henricksen, 2003; Hull, 1997; Kokinov, 1999; Ryan et al., 1997; Schilit & Theimer, 1994; Vieira et al., 2011). Many of these definitions are circular (i.e., use synonyms) or define context by example. Nevertheless, even among these weaker definitions, what is largely agreed upon is that context involves the answers to who, what, when, where, why, and how (Fischer, 2012; Jang & Woo, 2003), includes only a subset of the information available to an entity (Ali et al., 2010; Zimmermann et al., 2007), and is related to the situation (Dey & Abowd, 1999; Shishkov et al., 2018), activity (Dourish, 2004), or goals (Ali et al., 2010) of the entity. The most widely accepted definition states:

Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves. (Dey & Abowd, 1999)

Dey (2001) later expanded upon this definition to assert that the situation is “a description of the states of relevant entities.” By using the word “states,” Dey implies that the description is not stable. Indeed, changes in context both result from and cause changes in situation (Shishkov et al., 2018). This chicken-and-egg scenario presents a critical challenge for context-aware applications. If the situation is used to determine what information is and is not context, then the situation needs to be represented before context is determined. However, understanding the situation requires identifying relevant entities. How can the relevant entities be determined before identifying the situation when the situation is used to determine relevance?

The paradox has been avoided by having application developers determine what information is context and what the machine should do with that information (Dey & Abowd, 1999). Generally, this has meant that context-aware applications are only useful in a limited set of situations. For example, if a user is driving, traffic information, available in most map apps, is relevant to the situation, but if they stop driving, the map apps do not try to determine if the situation has changed and do not offer new functionality specific to the new situation (e.g., realize the car has a flat tire and present numbers of local roadside assistance companies). Instead, the application continues to provide traffic data. Even among the most sophisticated instantiations of computational communications agents, the applicable situations are limited (e.g., Myers et al., 2019). The problem is not trivial. Situations in the military can shift rapidly and as the situation shifts, the conversation shifts, and new, unexpected elements of the environment become relevant and older elements can decline in relevance. Thus, a model of context-awareness needs to account for a shifting landscape.

Zimmermann et al. (2007) elaborated on the representations of context by first identifying a core set of five context dimensions (individuality, activity, location, time, and relations) and their interactions (see Figure 1). Each of these dimensions corresponds to some aspect of who, what, when, where, why, or how. As with Dey & Abowd (1999), the context is centered on the entity, which can be animate or inanimate. Critically, the entity is the one being observed, and it may or may not be capable of representing context itself. Thus, in the 2014 incident involving friendly fire, one of the entities was the B-1. The context of B-1 was represented by multiple humans: Air Force control, JTAC, and the pilots. The B-1 had no ability to form context representation. Each of the humans was also an entity. JTAC’s context could be represented by the Air Force controller, the pilots, and, also, by the JTAC himself.




Figure 1. Zimmermann et al. (2007) Model of Context.

We can understand the role of Zimmermann et al.'s (2007) dimensions of context awareness best by stepping through an example using the B-1 as an entity. The B-1's individuality included *who* (or in this case *what*) it was. This information can be particular to a specific entity or to a classification of the entity. For example, the classification of the plane as a B-1 meant it had a 4-person team and could drop multiple bombs. Further, the B-1 did not have the ability to detect strobe lights usually used by ground troops to signal ally location. This particular B-1 also had a unique feature not available in all other B-1s: a precision targeting system.

The entity's location (*where*) is a second dimension of context. The location information can be relative to another entity (e.g., straight above the ground troops) or objective (e.g., precise global positioning system [GPS] coordinates). The third dimension, time (*when*) also offers relative and objective information. For the B-1, the interaction between location and time was particularly important in terms of how long it would take the B-1 to reach the troops.

Relations is a fourth dimension of context and includes *how* an entity interacts with other entities and itself. Relations can be social, functional, or compositional in nature. Social relations (e.g., rank) do not exist for non-agent entities. Thus, the Air Force controller's representation of the B-1's pilot included a social relation, but his representation of the B-1 did not. Functional relations include how the entity can use other available entities (or be used by them) to affect change. The B-1 had a functional relationship with the pilots, who were directly manipulating it. Zimmermann et al. (2007) take functional relations a step further to state that they also include how the entity affects thought. For example, the B-1's presence, initially, made the ground troops feel more secure and lessened their anxiety about the situation. Compositional relations relate to the component parts of an entity (e.g., landing gear). In a compositional relationship the parts

cease to function if the whole is destroyed but the whole *may* be able to function without some of its parts.

The final dimension is activity (*what* the entity is doing and *why* it is doing it). The B-1 had two core activities during the situation, fly and drop the bomb. As with social relationships, a representation for the goals of the B-1 cannot be directly drafted. However, the B-1 also interacts functionally with the pilots, whose goals can be represented. *Why* the B-1 is flying or dropping a bomb exists in its interaction between activity and functional relations.

Some of the context identified by Zimmermann et al. (2007) is detectable in the moment (e.g., *what* an entity is doing and *where* it is). However, other information cannot be detected and exists only as declarative knowledge or through inference (e.g., the goals of the entity). Goals are particularly important for understanding the situation as they make it possible to predict the entities actions. For example, understanding that the goal of the B-1 is to provide air support allows the prediction that it will drop bombs. In a different situation, an air show for example, the goal might be to highlight its speed and therefore a prediction about it dropping a bomb would not be consistent with its contextual representation. Detecting goals is difficult for both humans and machines, but unlike machines, humans can rely on a rich set of scripts developed throughout their lifetime to infer the goals and intentions of others. We know for example, that most people have an inherent goal to live and will take actions in support of that goal. We also know (or at least believe we know) why people are waiting at a bus stop, standing in front of an ATM, etc.

One of the strengths of dividing context into clear dimensions is that it allows an agent to represent the absence of information. The agent can then seek out or ask for the missing information. However, there will also be times when missing information is not relevant (Dey & Abowd, 1999). Thus, we circle back to the inherent difficulty in assessing relevance. We introduce three points about context representations which we argue will allow a machine to resolve this paradox: (a) they contain error, (b) they fluctuate, and (c) they require theory of mind. We elaborate on these points below to suggest principles which could be used to guide the development of context-aware, situation-independent machines.

2.1.1 Error in the Representation

What is critical about the representation of the entity is that it is created by an observing agent. Therefore, some information about the entity will not be available to the agent, leaving the representation incomplete. Further, because each agent has a unique perspective of the entity, no two representations of the entity will be the same (Zimmermann et al., 2007). For example, the B-1's location was unknown to the JTAC but known to the Air Force controller.

The representation may also be incorrect. All parties, except the pilot, represented the B-1 as capable of detecting strobe lights. Critically, this was false. This motivates the first principle guiding development of a context-aware application: *The application should iteratively update its model, seek out missing information, and be able to assess its confidence that represented information is correct.*

2.1.2 Fluctuations According to Attention and Focus

Each iteration should strive to prioritize the discovery and representation of information relevant in the moment. When the JTAC first made the request for air support, the B-1's exact location

was not as important as how quickly the B-1 could arrive. At the point the B-1 arrived, location increased in relevance as it was useful in determining why the communication was inconsistent (e.g., the B-1 was over a mountain range). Thus, the precise location information became relevant to the JTAC, and his representation of the B-1 shifted.

In Dey and Abowd's (1999) definition, all pieces of contextual information that characterize the situation are equal, but in the example above we can see that this is not the case. The location of the B-1 is not initially available or relevant but becomes relevant as the location begins to affect the situation. Even still, the location information is not as important as the activity of the B1. Thus, a problem with Dey's definition is that it only allows for a dichotomy: information characterizes the situation, or it does not. Zimmermann et al. (2007) contend that information relevance is based on its point of origin and activity. Contextual information decreases in relevance over time, i.e., it becomes outdated which can decrease the accuracy of the information or its applicability to the current moment. A representation of the B-1's location a moment ago is not as relevant as it was a moment ago. That is, the B-1 is in motion and its exact location is constantly in flux. Further, information connected to the activity (including goals of the entity) is likely to be more relevant than information not directly related to the activity. Thus, the functional relationship between the pilot and the B-1 is of high relevance because it is related to the activity (flying), and the compositional relationship of the landing gear to the B-1 is of low relevance because the landing gear is not needed to perform the activity of flying.

Another way of looking at this is that the applicability of the information reflects its relevance (Schmidt, 2002). Thus, relevance can be partially assessed by determining if the entity is using the information. Indeed, context has at times been defined simply as information *in use* (Dourish, 2004; Kayes et al., 2015; Winograd, 2001). For a non-autonomous entity, *in use* can relate to functional parts. The B-1 is using its engine and precision targeting system. For an autonomous entity, *in use* also reflects the attention of the entity.

Critically, information in use does not *necessarily* reflect an objective measure of the relevance or *usefulness* to the situation. Attention is error prone. For example, the JTAC's attention was on elimination of enemy combatants. This is not incorrect, but in attending so closely to this information, he lost track of the location of a unit separated from him. He then provided information that was out-of-date. Nevertheless, the information the JTAC is using is important to characterizing his situation. From this, we can draft an inference that his intelligence on enemy troops is likely to be more accurate than on allied troops.¹ The JTAC in this situation certainly did not intend to under attend to ally location. Attention is not fully governed by intentionality. It is also affected by salience, spreading activation, and familiarity. Certainly, as it posed a risk to his life, enemy troop location could be surmised to have been particularly salient to the JTAC in this situation.

From the role of origin, activity, and attention, we establish the second principle: *All available information about an entity exists on a gradient with the information that is actively in use by the entity as the most relevant characterization of it in the situation.*

¹ Communication transcripts are not available from this incident, but in some instances, the attention of the JTAC could be detected from patterns of speech, such as how often enemies versus allies are mentioned and what information is provided unprompted.

2.1.3 Theory of Mind

Context representations will also involve some theory of mind. An agent needs to know what information is available to an entity in addition to knowing what information is relevant to the entity. For example, if two planes are on a collision course with each other, air traffic control may recognize this information as characterizing the situation of the pilots. However, it will not help with predicting behavior or, critically, communication, unless the controller also recognizes that the pilots are unaware of their trajectory. How agents determine that information is available to an entity is thus of critical importance. Briefly, this can be accomplished by assessing the perceptual information currently available to the entity, knowledge about the entity (e.g., is it reasonable to assume the entity was co-present with the information; Clark & Marshall, 1981), and through communication. We will explore this further in Section 2.2 and Section 3. For now, we establish this as the final principle: *The agent must be able to form a representation of the context available to the entity and its ability to use that information.*

2.1.4 Situational versus Contextual Awareness

Having a representation of context is necessary but not sufficient for context awareness. To be aware the agent must be able to use information about an entity's situation in execution of a function (Dey & Abowd, 1999). The situation is the aggregate context of the relevant entities (Dey, 2001). Situation awareness can then be thought of as the degree to which an agent has accurately represented the aggregate context of the entities and can use this information to successfully carry out an action. This conceptualization of situation awareness fits with previously established definitions. Prior definitions characterize situation awareness as the understanding of all the events at play and how new events can emerge from prior events (Baumgartner et al., 2010; Endsley, 1995). Events can be understood as a segment of time observed by the agent to have a beginning and an end (Zacks & Tversky, 2001). The situation is thus the accumulation of the events currently perceived and the events which are expected to occur. The difference between how we have been discussing context and how situation awareness is discussed is that the context is centered on the entity. Ultimately, it is the context of the entities who engage in events. Thus, we can determine the relevant entities to a situation by determining what entities are involved in an event, though determining what events are relevant will still be a hurdle.

2.2 Defining Context-Aware *Communication*

If we define context as the information that “characterizes the situation of an entity” and awareness as the use of this information (Dey & Abowd, 1999), then we might define context-aware communication as an attempt to use context when updating one or more agents' characterization of the situation of one or more entities. The implication of “updating” a representation is that an initial representation must exist before the communication can begin. Thus, no communication can exist without context and a critical role for the producer² is to determine what already exists and needs to be added to the comprehender's representation of relevant entities.

² Throughout this section, we use producer to refer to the agent who is attempting to convey information and the comprehender as the person attempting to understand. During communication, the producer and comprehender roles will switch multiple times.

Contextual information during communication may come from (a) the environment, (b) memory, (c) the text³ itself, and (d) familiarity with the language including paralinguistic features. Each of these represents a distinct challenge for machine agents intending to communicate with humans. To understand these challenges, we first explore what we mean by a representation of the text and then how each type of context serves to facilitate communication.

2.2.1 Situation Model

Early research into how comprehenders represent language made it abundantly clear that comprehenders do not fully retain, even in the short-term, a representation of the exact lexical and syntactic values of the text (Bransford et al., 1972; Fletcher & Chrysler, 1990; Garnham, 1982; Glenberg et al., 1987; Kintsch et al., 1990; Radvansky et al., 1990; Schmalhofer & Glavanov, 1986; Zwaan & Radvansky, 1998). Instead, comprehenders integrate the text with their prior knowledge into a situation model (Kintsch & van Dijk, 1983). Ultimately, the situation model represents what the comprehender believes is important to understand and retain from the communication and thus also represents the text as filtered by the comprehender's goals (McCrudden & Schraw, 2007; van den Broek et al., 1995). Because the comprehender's situation model is, in part, a function of their individuality, the model will not precisely reflect the intentions of the producer. Thus, for the producer to be understood, they must consider the lens through which the comprehender will be processing the information. This use of theory of mind as part of conversation is why principle 3 from the previous section is critical to context-aware machine agents. Understanding the information needs of the comprehender is a critical task for the producer and one which even sophisticated machine agents struggle with (McNeese et al., 2018).

The producer also has a situation model for what they are about to communicate. Like the comprehender's model, it is ultimately richer than the lexical and syntactic values they will produce. Even if the producer were to attempt to articulate every part of the situation model, they cannot fully relay their perceptual and historical experience. Further, providing too much information risks the critical, new information being lost. Thus, the producer must discern how much of their model should be encoded into the text by assessing the context available to the comprehender. The context that will be omitted from the text is ideally shared between the producer and the comprehender and can therefore be inferred by the comprehender.

Communicators rely on a co-presence heuristic to determine if context is shared (Clark & Marshall, 1981). This heuristic suggests that information is shared if both communicators are aware that the information was at one point accessible to the other. When communicators are unsure of the accessibility of information, they may engage in grounding (Clark & Schaefer, 1987) during which they determine the level of context available to each party and provide new context critical to comprehending the producer's situation model. During grounding, a communicator might give positive (nodding) or negative evidence (request for clarification) that enough information is available to understand the communication (Clark & Brennan, 1991).

Using concepts of co-presence and grounding, we can now demonstrate how the environment, memory, text, and familiarity with the language contribute to the formation of a situation model.

³ We use text here to refer to communication across modalities, not just written communication.

2.2.1.1 Environmental Context

The environment refers to the perceptual experience of the communicators. This experience will be informed partly by knowledge. For example, if a student is in an office with their professor, the individuality of the professor is relevant context that informs the student's perceptions. In a toy example, we will say this is the student's professor of colonial history and this is her office hour dedicated to meeting with students from the class. In (1) the professor is communicating a point to the student. The professor draws upon her perception of the student's individuality (accepted to college, new to the topic) and prior activity (attended class) in constructing the text. The professor thus assumes information stated in lecture and a basic understanding of France and Britain as one-time colonial powers is co-present. However, she will not assume the student has the same degree of knowledge as her colleagues.

- (1) The British and the French scientists were searching for the red-shanked douc in Indonesia. The BRITISH found the monkey in Indonesia.

Similarly, in comprehending this sentence, the student draws on their environment to understand that the goal of the communication is to learn something about colonialism rather than, say, conservation. Thus, the professor has drafted a simplified rendition of her situation model, and the student, in assembling their own model of the text, may focus more on the countries than the name of the specific type of monkey.

2.2.1.2 Memory Context

Continuing our example, each agent also uses their memory to inform their context. For both the professor and student, there are rich representations of learned information about the mentioned entities. The professor's understanding of what she is saying may be richer than the student's because she has more knowledge on the topic. However, for the student to represent and remember what the professor is saying, they will need to access the declarative knowledge of the topic that they do have and update that knowledge with the new information.

2.2.1.3 Textual Context

Context can also be found within the text. The student initially may not know what a red-shanked douc is, but later context from the text makes it clear that it is a monkey. Further, repetition of language across sentences allows the student to understand that the people who found the monkey were scientists, even though the term 'scientists' is not reproduced.

2.2.1.4 Language Familiarity Context

Finally, familiarity with the language is another form of context used in this example. At its most basic, both the professor and student know English and will use this knowledge in drafting their communications (e.g., speed and complexity of the language used). Further, even if they are not aware of it, the familiarity will increase how much information they can derive from the text. For example, in the second sentence of (1), *British* is capitalized to indicate a L + H* prosodic structure⁴ on the word. This pitch accent indicates contrast (Zimmermann, 2008). In other words, the pitch accent on *British* cues the comprehender to contrast *British* with *French* (Fraundorf et

⁴ L + H* is used in the ToBI system to denote a contrastive pitch accent in which the pitch transitions from low to high with stress placed on the higher pitch.

al., 2010). The comprehender thus uses their experience with English as context to know that not only did the *British* find the monkey, the *French* did not, *and* the fact that the French did not is important. Indeed, comprehenders are more likely to remember that the *French* did not find the monkey when an L + H* pitch accent is placed on *British* or, in writing, when *British* is capitalized, italicized, or preceded by a cleft (e.g., *it was the British*; Fraundorf et al., 2013; Norberg & Fraundorf, 2021).

Ultimately, the cumulative context facilitates integration of the text so that what the student constructs as their situation model and remembers may be akin to (2). The addition of *colonial activities* in (2) is a critical part of understanding the concept of the situation model. Although the professor never mentioned colonialism in the utterance, the context triggered the student to update their understanding in relation to colonialism.

- (2) The British and French scientists were carrying out colonial activities in Indonesia and looking for a monkey. The British found the monkey and the French did not.

2.2.2 Grounding

A producer does not want to say something that cannot be understood because the comprehender does not have the correct context. That is, the professor did not want their student to go home and study red shanked doucs and ignore French and British rivalries. However, the professor believed that the student had enough context to understand the goals of the statement. Similarly, a comprehender does not want to misunderstand what the producer means because they applied the wrong context. Communicators use grounding to help avoid errors related to contextual accessibility. A full overview of grounding is outside the scope of this review, but we note that grounding is concerned with ensuring that both parties have the information they need to understand a communication. For example, if the professor was not confident in their memory of seeing the student in class, she might have begun by asking, “Were you in class earlier today?” to ensure that the student had the requisite knowledge from that day’s lecture. Similarly, the student might have expressed that the information was at too high a level by furrowing their brow to alert the professor that they were missing context.

Critically, grounding is not just ensuring that information is available, it is ensuring it is accessible. It is at times necessary to provide mutually known information in a communication to ensure that the information is *in use* by the comprehender. This is because having information in memory is not the same as being able to retrieve it.

One of the most common constructions in an ongoing communication is the given-new structure in which the producer first provides information that the comprehender *already* knows and then provides new information. The producer can alert the comprehender about which information is meant to be given (contextual) with a H* pitch accent and information that is not shared with a L* pitch accent (Schwarzchild, 1999; Zimmerman, 2008). The subtle difference in accenting provides information to the comprehender about what the producer *believes* is co-present knowledge. For example, in (3) the producer believes the comprehender knows which trousers they are referring to but is uncertain about whether the comprehender really put them in the microwave and thus does not express this information with the H* accent.

- (3) You put my TROUSERS in the microwave? (All caps represent a H* pitch accent)

Providing mutually known information can serve two purposes: It can ensure the comprehender is using the information and it connects information across multiple exchanges. Thus, part of grounding is the producer ensuring that co-present information is in fact in use. Detecting given-new patterns in communication could help a machine agent determine what is meant to be already known and what is meant to be new information to add to its database. An agent which can detect prosody might be trained to detect these prosodic features, but prosody alone is likely not reliable enough. There are also indications that givenness can be marked by a definite article (Haviland & Clark, 1974; Kintsch, 1998; Yang et al., 2007) and predictive models using latent semantic analysis have established a way to determine whether information in text is meant to be given or new at 74-80% accuracy (Hempelmann et al., 2005).

These types of distinctions could ultimately help a machine agent determine which information in an utterance should cue it to search for shared context and which should cue it to add new information. Further, the machine agent could use this to ensure that humans are accurate in their assumptions about shared context. If they indicated that information about the trousers should be known to all parties but the machine agent can detect that one party had stepped away from the communication when *trousers* were first mentioned, an alert could be sent to provide more information about the trousers.

2.2.3 Miscommunication and Clarifications

As we have illustrated, communication hinges on successful use of context by both the comprehender and the producer and specifically on identifying what context the other is currently using or has readily available. This will always be, at least in part, an inference, and misconceptions about what is co-present can lead to communication challenges. One of the underlying ways errors in co-presence can occur is the assumption that because information is available to one, it is available to all. This ego-centric view can lead to a producer providing too little context or to a comprehender assigning information to an entity unavailable to the producer. For example, a producer might say, “He’s going to get hurt,” assuming that information about *hurt* will draw the comprehender’s attention to something that was said earlier. However, the comprehender may not be thinking of the earlier conversation and instead may apply it to a person they are looking at outside the window, even though the comprehender *knows* the producer is not looking out the window. Paradoxically, this type of ego-centric failure of context awareness typically increases with familiarity (Mustajoki, 2012; Savitsky et al., 2010).

In the above example, the communication error was unrecognized. However, even if the comprehender asks, “who is *he*?” time has been lost, and in operational communication, time is critical. Indeed, late information is a critical issue in communications during search and rescue training missions (Rothwell et al., 2022). Clarifications, while better than misunderstanding, risk decreasing timeliness.

Nevertheless, a critical question regarding communication is why do comprehenders sometimes ask for clarification and sometimes assume understanding. The recipe for failing to recognize that comprehension has failed is no doubt multi-layered. It must first involve the possibility of comprehension. If the comprehender is completely unable to integrate what was said with their prior knowledge, a situation model cannot be drafted and the comprehender will be alerted to an error. Indeed, comprehenders are more likely to form misunderstandings when a partial situation model could be created and the process of forming that situation model *felt* easy (Mata, 2020;

Norberg, 2022). This has its origins in the concept of good-enough processing which theorizes that comprehenders process communication with only enough depth to get a gist (Ferreira & Patson, 2007). If a gist is quickly conceived, then the comprehender is less likely to interrogate their understanding. Thus, miscommunication might be most likely when the comprehender is *most familiar* with the topic or situation. Indeed, this would account for findings that familiarity increases ego-centric uses of context.

In the bombing incident, the pilots were given conflicting information about the location of the allied troops and enemies. It was not that the pilots were told they were in one location and in the next the location shifted. The pilots were told troop locations but the information they received next about gun fire directionality was inconsistent with these locations. The pilots did not attempt to integrate the information and instead, integrated the troop location and gun fire directionality into a situation model but not with each other. It is difficult to say why the pilots did not notice the discrepancy. It may be that the scenario was familiar enough to them that they believed they had a fully represented model and so did not interrogate their model any further. This is precisely the type of scenario during which a machine agent who is listening to the conversation may be able to intervene. The machine agent is less likely to take the short-cuts humans make in their processing and can therefore more readily detect when communications do not match.

2.2.4 Team Communication

While attributes of context-aware communication apply to teams, we have primarily been discussing one-to-one communication processes. Team communication is paramount in the military and must also be considered. In the military, in addition to one-to-one communications, there are one-to-many (commander to troops), many-to-one (troops to command), and many-to-many (everyone talking to each other) communication. Each of these has slightly different concerns regarding contextual accessibility and co-presence. In a one-to-many context, the one needs to track comprehension across multiple parties simultaneously and make decisions about whether to form communications based on information co-present to all (a lowest common denominator approach), information co-present to most, or to direct statements to specific agents within the larger group (e.g., Squadron 52).

During many-to-one communication, the information only needs to be co-present between the producer and the one comprehender, but the producer must consider that this information will not have the same rate of accessibility to the comprehender because the comprehender is attempting to keep multiple strains of communication active. The producer will thus need to do more work to ensure grounding. Further, because the comprehender can only process information from one producer at a time, the producer will need to determine the priority of their information in light of all other information and take care to avoid speaking too long so other communications can get through. Thus, the producer must be careful to give enough information but also ensure that passing the information does not take longer than absolutely necessary. Concerns about quantity and manner have been central to discussions of all communication (Grice, 1975), and can be thought of as a desire to minimize collaborative effort (Clark & Brennan, 1991). However, collaboration takes on new constraints when the focus of the collaborator is divided.

Finally, many-to-many communication requires the combination of considerations for many-to-one and one-to-many. The producer must still ensure that they do not speak for longer than necessary and must also decide whether they will rely on co-present context available to some or all of the comprehenders. Finally, the comprehenders will have the added task of determining which parts of the communication are relevant to their situation.

Team structure can also be a critical part of the context of communication. The structure of the team and the workload may require different patterns of communication (Marlow et al., 2018). For example, hierarchical teams, like those in the military, function better with less frequent communication (Urban et al., 1995; Kleinman & Serfaty, 1989). Hierarchical structures may be able to implicitly coordinate due to a deep understanding of each other's roles (Volpe et al., 1996). Thus, when a situation changes, each member already knows how everyone will respond and trusts everyone to respond as expected, and therefore no communication is necessary (McChrystal et al., 2015).

High workload environments are also more successful when the frequency of communication is lower (though it is the opposite for low and moderate workload environments). It is unclear why less communication is necessary (or useful) within high workload environments. It may be that high workload environments increase monitoring of other teammates as the activity of one member is more critical to the moment (Marlow et al., 2018). On the other hand, in a lower workload context the shared model can develop more gradually (Cannon-Bowers et al., 1993).

In considering a context-aware computational communications agent, it will be necessary to consider the utility and application of the device across team contexts. The agent will need to consider the nature of its communication (how many listeners) in order to adapt its communication. Further, given that military operations are often low frequency communications environments, a communications agent may need to be given access to higher frequency communications in order to establish critical elements of the situation (e.g., relevant entities and goals).

3.0 THE ABC MODEL OF CONTEXT-AWARE COMMUNICATION

We draw heavily on the work of Zimmermann et al. (2007), Dey and Abowd (1999), and Clark and Marshall (1981) in drafting a new model of context-aware communication. The goal of the model is two-fold. First, prior models of context-awareness do not represent how co-present information is used and created nor do they provide a path to considering when co-present information needs to be grounded (stated explicitly despite being known to all communicators). In other words, communicators are ABC, and our model reflects this process. The second goal is to consider how a machine agent might adapt its functionality across disparate situations. Critically, the machine agent posited here could support modules which feed information from the environment and sort it according to its dimension of context. Thus, in one instantiation, the machine agent might receive location information from cell phone towers but in another it would receive this information from radar. The information from both could be “plugged into” the agents’ processor for representing location. Thus, in the model, we are agnostic about the mechanism for receiving information outside of the communication as the mechanism will be domain specific.

Our model (Figure 2) begins with two communicators (Agent 1 and Agent 2); however, we note that the principles apply to multiple communicators. Further, the agents may be human or machine. We first consider that the conversation is happening as part of a situation (grey circle). Second, each communicator has information that only they know. In any communication, at least one of the agents has chosen to relay some subset of this information. By virtue of being stated, this information is necessarily *in use* by the producer and is therefore relevant to the characterization of the situation of the producer. *If* the comprehender is attending to the communication, it is also in use by the comprehender. We further place a spillover zone around the communicated context to represent the information that the communicators have active due to the conversation, but which is not shared in the communication (pink glow surrounding communicated context in use in Figure 2).




Figure 2. ABC Model of Context-Aware Communication.

Note. The abbreviation 'Comm' in the center of the situation model stands for 'Communicated'.

Continuing through the model, once information about an entity has been communicated, it becomes co-present. All agents in the conversation now consider this information to be mutually known. Co-present context includes all the information that the agents both know regardless of whether it was recently communicated. Agents can use the differentiation between what is co-present and what only they know to guide production and comprehension processes.

In our model, we save co-presence for information that truly was made available to both agents through explicit communication. If Agent 1 believes Agent 2 knows something, but that belief is based on an inference regarding co-presence, then it is represented as part of Agent 1's personal context. This is because Agent 2 has no way of knowing what information Agent 1 has inferred, so it cannot be co-present. For example, the pilot and the JTAC knew that the pilots were flying a B-1. That was co-present. It was stated in the conversation. However, the pilots assumed that details about the B-1, like that it could not detect strobe lights, were known to the JTAC. But the pilots actually had no specific reason to infer that this information was ever provided to the controller, so they made an inference about its availability to JTAC.

Critically, information in each segment of the model is not equally available to the agents. To represent this, information is represented on a gradient (changes in color hue in Figure 2 triangles). At the start of the conversation, the gradient of available information will be influenced by a combination of information salience, frequency of accessing a particular piece of information, and the current goals of the communicator. However, as the conversation continues, the gradient will be shaped via spreading activation from the content of the conversation (communicated context in use), making information relevant to the conversation more accessible and speeding integration processes when creating the situation model.

Due to cognitive capacity constraints and time constraints, the information in use during a conversation is going to be limited. Thus, even though this information is co-present, its accessibility to each agent will decay over time. Time since last access is thus another variable affecting the gradient. If information has not been mentioned in a while, it will need to be restated (i.e., grounded). Finally, the co-presence gradients will be separate for each communicator, but critically, the degree of availability itself is not co-present. Agent 1 cannot be certain that information that is co-present is as readily accessible to Agent 2 as it is to themselves.

This model provides a starting place for the representation of context. However, for an agent to be context-aware, they must be able to use this information, either to engage in the conversation or perform an action in the environment. In the following section we highlight how this model supports awareness and provides a machine agent a basis from which to engage in context-aware human-human and human-machine communication.

4.0 BUILDING A CONTEXT-AWARE COMMUNICATIONS AGENT

The conceptual model in Figure 2 illustrates how a machine agent might accomplish the principles outlined in Section 2.1: (a) The model is updatable. (b) A gradient of relevance is established. (c) It considers the perspective of the other agent by separating private and co-present knowledge, by making inferences about the knowledge of the other agent, and by drafting a representation of how likely that knowledge is to be readily accessible to the other agent.

There are also material needs that are implicated by the model. The machine agent must have a data frame of networked information that is updatable and retrievable, and the networked data must contain connections to domains, other information, and broader situations. The data frame (or at least part of it) must be structured around entities which the agent must be able to identify in the environment. Finally, the machine agent must be able to act. It must *use* this information appropriately, which Fischer (2012) defined as getting the right information, at the right time, to the right place, via the right way (modality), and to the right person.

Realization of this model will require advancement in multiple technological areas. We focus here on the areas related to using rather than detecting context. Current context-aware communication systems broadly fall into one of four categories proposed by Pascoe (1997): sensing, augmenting, resource exploitation, and adaptation. These categories were devised early in the creation of context-aware computing but remain broadly applicable. According to Pascoe, *sensing* devices can detect information in the environment and supply it to the user whereas *augmenting* devices have the added ability of associating information with a specific context (Dey & Abowd, 1999). *Adapting* devices change their functioning based on the context, and *resource exploiting* devices network with other computing devices to accomplish goals (either sharing information or facilitating the commands of the user).

In our discussion of modern context-aware applications and proposals for context-aware communications in enhanced operations, we reinterpret the categories to reflect the technology today more accurately. First, sensing underpins all the other capabilities. Although Pascoe (1997) stated that sensing devices also supplied information, we categorize all information supplying devices under the category of augmentation and stipulate that sensing is a prerequisite ability for any context-aware computing. Second, Pascoe restricted adaptation to the machine's ability to adapt itself. However, we extend this ability to the machine's ability to adapt the situation.

4.1 Context Sensing as a Prerequisite

Devices which can sense context are widespread in contemporary society. All applications which detect the user's location or draw on specific information in a user profile are context sensing. Many of today's context-sensing apps and devices focus on a subset of Zimmermann's (2007) dimensions of context. They detect location and individuality information (e.g., online search histories for targeted advertisements) or they detect information about time, location, and activity (e.g., Fitbit). However, they rarely create a full contextual representation of all five of Zimmermann's dimensions and their interactions. Further, they typically draft only a contextual representation of a single user. The devices rarely attempt to sense the context of other entities in the user's environment.

Context sensing is a critical part of the development of any context-aware application because the machine agent will not be able to sense everything. Developers must decide what information

the agent will need to carry out its task. As a result, the sensors provided to machine agents are highly situation dependent. What a machine agent needs to operate in search and rescue missions is quite different from what it needs to assist in a response to an active shooter or to mission planning. In considering the information needs of a context-sensing communications machine agent we emphasize a focus on the operational requirements. If it is to engage in communication, it must have access to the communication. All other sensors should be modular, that is able to interface with the core functioning. Thus, the core system should have an ability to represent location, individuality, time, relations, and activity from the conversation. Then, if it is operating in an aviation environment, radar can be “plugged into” the location representation to provide additional location context it can use to understand the communication or provide information.

The context-aware application can then use the communicated context-in-use and goals (e.g., from listening to a briefing) to inform the gradient and prioritize data collected from its sensors. Thus, information about strobe-lights becomes important if referenced in a conversation. Further, if it has the ability to detect enemy planes and the briefing indicated that this was a priority, it will continue to monitor for enemy planes despite lack of recent chatter on the topic on the radio.

Briefly, we believe that the representations can be drafted by training a natural language processing model on text annotated for named-entity recognition, part-of-speech, and context dimensions. Access to a corpus and use of an ontology and formal semantic representation which can be networked will also be necessary, similar to Ball et al. (2009). Further, representations from multiple users might be developed by taking advantage of an internet of things. That is, an application can exist on multiple devices drafting representations of single users, and a centralized application may attempt to aggregate the information, thus forming additional entity representations.

4.2 Context Augmenting Technologies

Context-aware computing devices require sensing. However, sensing is not awareness. Ultimately, it is necessary for the machine agent to use the information it has sensed and represented. In augmentation, this is in the form of providing communication, i.e., identifying what information a second agent needs and providing it to that agent.

There are many devices that currently have this type of functionality although most of them engage only in one-way communications meant to inform or alert the user to context. Further, these applications may be designed to provide information to facilitate a user’s context awareness, *or* the applications may be situated between two users to facilitate communication. For example, livestream traffic data and product recommenders (e.g., online ads, movie recommendations) typically only include, or at least they assume, a single user interacting with the application. The goal is to make the user aware of information that may be relevant to their situation. On the other hand, applications which tell users about the availability status (e.g., away, busy, idle) of another user or provide information about the sentiment conveyed in an email (e.g., Grammarly) exist between two users and are aimed at facilitating the communication.

Context augmentation can also go beyond sharing context. Augmentation applications may use context to determine *when* and *how* information should be revealed. For example, alerts have been designed to trigger based on location (e.g., Amber alerts), time (e.g., flight status), detection of activity in the environment (e.g., step reminders in health apps), relative location to another

user (e.g., Lovegely; Iwatani, 1998), or any of those combinations (Dey et al., 2001; Fogarty et al., 2004; Marmasse & Schmandt, 2000; Miluzzo et al., 2008; Mynatt et al., 1998; Ranganathan & Lei, 2003; Siewiorek et al., 2005). One early application was particularly applicable to enhanced operational communications. Secure Integrated Response Electronic Notification (SIREN) (Jiang et al., 2004) detected activity and location to manage communications among firefighters. When it detected a hazard, the application would alert all firefighters to the problem. For example, it would automatically sense the heat (activity) coming off a room, alert nearby firefighters, and mark that room as unavailable during an egress.

In each of these examples, the augmentation is in the form of an alert or short, one-way communication. There are fewer computational agents who are designed for both conversation and context-awareness. ConChat was one such application. It allowed users to make information about their location, room temperature, noise, etc. available to other users and to the application itself (Ranganathan & Lei, 2003). ConChat then monitored conversations to detect circumstances when divergences in context relevant to the conversation existed between two users. For example, if users were in different time zones and one user said “2:00PM,” the application would augment the original provided time with the other users’ relative time.

Other modern uses of conversational context augmentation target specific subsets of people. For example, some modern devices are seeking to speed up communication with patients who have aphasia or who are paralyzed by utilizing context to help the patient locate appropriate words (Davis et al., 2003; Kane et al., 2012). Although predictive text also exists in many applications today, it is not sufficient for the needs of some disabled users. Thus, a narrower, context-appropriate set of suggested vocabulary greatly speeds their ability to communicate while lessening the cognitive load involved in producing an utterance. These devices create profiles of people and locations with which the patient commonly interacts. The application can then use facial recognition and GPS coordinates to determine contextually relevant vocabulary. These are some of the rare devices which seek to identify and represent relevant entities beyond the user.

All the above devices move the private context of the machine into co-present context by providing human agents with information they didn’t otherwise have access to. Availability of this information augments the human agent’s ability to understand others and the situation. They also provide examples of how context from the environment (e.g., SIREN), the text interacting with the environment (e.g., ConChat), or the language itself (e.g., Grammarly) might be used to inform context representations. Memory is even being used as context now in Microsoft Exchange where an out-of-office auto reply is only provided once when the first email is sent to an away user rather than providing the auto reply to every message sent to the user while they are away. However, none of these devices use all four sources of context at once, and most critically, none of them work in novel situations. They cannot selectively determine across multiple sensors what information is relevant. We believe our model of context will allow for a more versatile machine agent. Now we highlight several possible context-aware augmentation applications that may have prevented the communications issues involving friendly fire in 2014.

The incident report criticized communication from the point when JTAC first radioed for air support until the moment the B-1 released the bomb (Lamothe, 2014).⁵ The report notes JTAC overstated the danger of the situation and the Air Force controller reported that he could “hear the stress in his voice” (Lamothe, 2014). It would have been difficult for a machine agent (or human) to know at the point of transmission if the need for air support was warranted. Indeed, the controller believed the stress the JTAC was clearly feeling was a signal *not* to question the JTAC despite misgivings about the focus of the JTAC’s attention.

However, an application connected to the radio (but using a data link) could have detected sentiment in the prosody and lexical choices of the JTAC and passed that information to the Air Force controller. With this individuality information in hand, the application could have prompted the controller to ask additional questions to ensure accuracy in the details. It is possible the controller would have been more comfortable pushing for more details if the advice were coming from a second source; however, this type of trust in machine teammate would need to be tested.

Ground forces could also have been equipped with a location technology that provided their relative location to each other. The application could have then checked the detected relative locations against any communications related to location. If the application detected an inconsistency, it could send an alert. In more general terms, the application would use the communicated context-in-use to assess what the human agents believe to be co-present. If it does not align with the application’s private context representation, it will know it needs to interrupt and share the discrepancy.

Further, if it could not detect the location, it could potentially detect that the JTAC has reason to be uncertain about the location of allied troops because his last confirmed awareness of their location was several minutes prior. This could have been used to trigger a request to radio and verify the location of ground units (if possible) or on a simpler scale provide a warning that the location data is out-of-date.

When the pilots received conflicting information about the location of friendly units, the application would not have even needed to detect information in the environment to alert the human agents to an issue. A listening agent could have detected the discrepancy and pushed an alert to question if the newest information was an update or correction of a prior statement. The key is that in creating a situation model for the ongoing text, the machine agent will be able to support the human agents using its greater processing power. As information slips down the gradient of co-presence accessibility for humans, it can also decay such that when it is recalled, the recollection is incomplete or there is interference from new information. This may cause humans to let contradictions go unnoticed. A machine agent could detect the likelihood this information was forgotten and act to remind the human agents (engage in grounding). Further, the information the machine agent recalls would not be subject to the same decay and interference.

Finally, and potentially most critically, the Air Force controller, JTAC, and ground units all believed the B-1 could detect strobe lights, and the ground units used strobe lights to identify

⁵ There are alternative accounts of this night, notably by the officers involved (Whitaker, 2017). We use this incident as an example but recognize that the reality was likely more complicated than outlined in news reports of the incident report.

their position that night. In fact, according to the controller, there was an unusual focus on the strobe lights by the ground troops (Lamothe, 2014). However, the B-1s did not have this capability. At one point, the Air Force controller asked the pilots if they could see the strobe lights and they responded “negative,” but the controller appears to have misunderstood the response. Both the controller and JTAC appeared to believe this meant the allies were clear of the zone where the B-1 would drop the bomb. They believed the transmission meant the pilots did not see strobe lights in the drop zone, but the pilots meant they had no way of seeing the strobe lights. A machine agent could have used its declarative knowledge to know the capabilities of the B-1. Access to the knowledge would have been cued through spreading activation by the repeated mention of strobe light detection in conjunction with a B-1. The conflict between using information about strobe lights when the entity which would need to detect the lights does not have this capability could have triggered an alert.

Finally, a truly sophisticated instantiation of the model would be able to serve in conversational roles traditionally requiring the cognitive and linguistic skills of people. In order to make this replacement, it would need an ability to represent goals and assess risk. The goal representations could be achieved by providing the agent information about mission objectives and linking that with sensor data. Risk assessment could be accomplished in part through evaluation of how events on the ground are affecting goals and in part by utilizing existing algorithms for assessing risk (e.g., Buyurgan & Lehlou, 2015).

There are several additional ways such an agent might function that do not fit the friendly fire incident, but which reflect known communication issues (e.g., Rothwell et al, 2022). First, missed communications can result when a human agent is navigating multiple communication channels. Thus, if Agent 1 opens a line of a communication with Agent 2, the machine agent could relay the status of Agent 2 including which channels Agent 2 is currently monitoring. This representation would come from a representation of what Agent 1 appears to believe about Agent 2 and by checking that against its own model of Agent 2 (e.g., currently talking on a separate line). Further, the machine agent could intervene and ask Agent 1 for the content of the intended transmission. By matching that against its current gradient of relevance regarding the goals and broader communications, it could infer if the information Agent 1 needs to share necessitates interrupting Agent 2.

Although many of the above examples come from a single situation, the agent needs to be able to adapt to new situations. Therefore, the central goal of the agent is to increase the probability of mission success by checking environment data against communication rather than by just providing environment data. For example, detecting relative location information of ground units down to the meter is something that could be determined a priori as likely relevant and useful information for commanders. However, before the soldiers hit by friendly fire entered the situation, no one knew they would need air support or that the support would come from a B-1, though it was a possibility. Thus, it would have been hard to realize before the incident that lack of prior knowledge regarding the B-1’s strobe light detection ability was going to instigate a fatal miscommunication. However, a machine agent which takes cues from the conversation and has a database filled with possibly useful information could be used in this instance to correct the miscommunication provided that it understood what information was and was not co-present to the entities in the conversation.

It must be noted though that detecting relevance via the communication has its own downsides. Namely, without a communication, the machine agent will have no way to determine the relevance of information. We believe this can be circumvented by ensuring access to mission plan content, so the machine at least has that plan representation as a starting point to use in assessing the context. A further way to increase the usefulness of the machine agent during times when communication is less frequent is to instill a capability for event representation into the agent. Thus, a review of detecting and representing events has been identified as a future research direction.

4.3 Context Adapting Technologies

Although context augmentation is the most common type of context-aware application and the one for which the range of use-cases is greatest, context-aware applications do not have to function as communicators themselves. Instead, they can use their representations to adapt the environment (e.g., computers changing display lighting at night, which makes reading easier and improves sleep).

An issue with the B-1 incident was that the signal-to-noise ratio in the radio communication was low, particularly when the B-1 flew over a nearby mountain range. A context adaptation mechanism that could either detect the quality of the actual connection or detect that there were too many requests for clarification or repair statements, then act to improve communication in two ways: (a) It could switch the communications to text using automatic speech recognition. (b) It could wait to send the transmission until connection bandwidth was high enough to prevent package loss. In both cases, the machine agent has detected context from either the text itself or the environment to determine that signal-to-noise is a problem and then acted upon the environment to address that issue.

The machine agent could also be taught to flag language patterns that can cause confusion (e.g., double negatives or using *not* in written text where it is often missed or even detecting if *not* was accidentally omitted by a producer during written communication). This type of capability already exists in applications such as Grammarly and could be adapted to fit the unique communication patterns in operational communication.

4.4 Context Resource Exploiting Technologies

Contextual resource exploitation improves the user's experience and facilitates their goals based on what resources the machine agent can exploit in the environment. For example, iPhones asking other iPhones to share Wi-Fi passwords is an example of a machine agent seeking to fulfill the goals of the user (logging onto Wi-Fi) by automatically searching for another machine that can provide the needed information.

A machine agent's ability to exploit resources is core to the proposed model. The model proposes an architecture that can adapt to the information it is fed. It exploits pre-existing sensor data. One simple use-case for exploitation is in detecting the availability of a controller. If Agent 1 attempts to contact Agent 2 and Agent 2 is busy, the machine agent could reroute the communication to another agent.

It also could exploit databases. If it needs to carry out a novel functionality, e.g., if ground units were diverted to an unexpected location and an extraction needs to be coordinated, it could use access to maps of the new location and active flight plans to help coordinate the extraction. Of

course, it will need some way of determining what information is needed, so there is again a core limitation to the functionality of the model in that it requires some form of communication to determine what types of information are relevant. However, simply asking the application to find information relevant to extracting units from a specific location at a specific time would potentially be enough. Further, if the computational agent is working with the ground units it could detect divergence of the current location from the planned location and recognize that terrain maps or other location information from its database will be relevant.

4.5 Related work at Air Force Research Laboratory (AFRL) - The Synthetic Teammate Project

There are several similarities between the model proposed here and a synthetic teammate that has been under development for the last 15 years or so at AFRL (Ball et al., 2009; Demir et al., 2015; Freiman et al., 2018; McNeese et al., 2018; Myers et al., 2019). The teammate specifically acts as an uninhabited air vehicle pilot for reconnaissance training missions.

Using Adaptive Control of Thought-Rational (ACT-R) (Anderson, 2007), researchers have developed a way for the synthetic teammate to merge information from the environment and the conversation. As in our proposed model, it does this by creating a situation model of the text. The result has been a machine teammate who is as successful during training missions as untrained humans (Myers et al., 2019). As sophisticated as this technology is, its creators have noted a few weak points (McNeese et al., 2018): (a) When there is no communication, the machine cycles through sensor data at random. (b) It has difficulty anticipating the information needs of its team. (c) It only works in this specific reconnaissance mission framework.

Our generalized model focuses more on the communication than the ability to carry out individual actions in the environment. This avoids the need for it determine how to fly, drive, swim, etc. and puts the focus on ensuring everyone is communicating effectively. An agent built from our conceptual model will not cycle through information randomly but will determine relevance of information over time.

By modelling the activity of team members, our model can also use spreading activation based on detection of activity to determine what information in its databases or coming from its sensors may be relevant to the humans. If the machine agent is given information about the mission, then there should be an overlap between the representation of one activity and a representation of the activity that will follow it. The machine can use the spreading activation from activity to activity to determine what information might be needed next. Critically, this will require the agent to predict what should happen next, what information is needed to complete that activity, and what information is *not* co-present.

The goals of the context-aware machine agent and the synthetic teammate under development are slightly different. The synthetic teammate communicates and performs activities to fulfill a specific mission agenda. To increase its generalizability, our machine agent will be designed first for communication alone. Nevertheless, we do not believe our model precludes an agent who can also pilot a Unmanned Aerial Vehicle (UAV). Rather, this type of functionality could be added based on domain-specific demands.

5.0 CONCLUSION

The definitions and models presented here allow us to move towards a dynamic context-aware agent which utilizes communication to understand the relative importance of information in the world (i.e., to recognize context). The machine agent's representation of context can then be leveraged to facilitate human-to-human communication as well as human-machine communication. We note this involves overcoming significant hurdles in NLP both in terms of the core functioning of the device and in terms of its application. For example, limitations in automatic speech recognition make it difficult for machine agents to get an accurate picture of the situation from speech alone. However, this technology is continuing to develop. As it does, we need to be prepared to implement new ways to use the information. This will include a more sophisticated way of handling incoming information.

While connecting speech to the environment is half of context, the other half is detecting the context of the text and familiarity with paralinguistic features. At the moment, NLP has difficulty utilizing information like the pace of the speech and detecting revisions to speech. It also has difficulty linking information across information exchanges. We believe the use of gradients for activation of information previously mentioned in speech will help with this but recognize these represent significant hurdles that researchers have been striving to address for years. Nevertheless, the architecture we propose here for representing context dimensions based on their shared overlap with other agents, and prioritizing context on a gradient could make significant steps forward that can have far reaching applications.

6.0 REFERENCES

- Ali, R., Dalpiaz, F., & Giorgini, P. (2010). A Goal-Based Framework for Contextual Requirements Modeling and Analysis. *Requirements Engineering*, 15(4), 439-458.
- Anderson, J. R. (1993). Production Systems and the ACT-R Theory. *Rules of the Mind*, 17-44.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.
- Badrinath, S., & Balakrishnan, H. (2022). Automatic Speech Recognition for Air Traffic Control Communications. *Transportation Research Record*, 2676(1), 798-810.
- Baddour, A. M., Sang, J., Hu, H., Akbar, M. A., Loulou, H., Ali, A., & Gulzar, K. (2019). CIM-CSS: A Formal Modeling Approach to Context Identification and Management for Intelligent Context-Sensitive Systems. *IEEE Access*, 7, 116056-116077.a
- Ball, J., Myers, C., Heiberg, A., Cooke, N. J., Matessa, M., Freiman, M., & Rodgers, S. (2010). The Synthetic Teammate Project. *Computational and Mathematical Organization Theory*, 16(3), 271-299.
- Baumgartner, N., Gottesheim, W., Mitsch, S., Retschitzegger, W., & Schwinger, W. (2010). BeAware!—Situation Awareness, The Ontology-Driven Way. *Data & Knowledge Engineering*, 69(11), 1181-1193.
- Bazire, M., & Brézillon, P. (2005, July). Understanding Context Before Using It. In *International and Interdisciplinary Conference on Modeling and Using Context* (pp. 29-40). Springer, Berlin, Heidelberg.
- Beck, L. J., & Cohen, H. H. (1994). Miscommunication and Human Error: The Difference Between Expectation and Reality is the Result of Communication Failure. *Ergonomics in Design*, 2(1), 16-20.
- Bransford, J. D., Barclay, J. R., & Franks, J. J. (1972). Sentence Memory: A Constructive versus Interpretive Approach. *Cognitive psychology*, 3(2), 193-209.
- Brézillon, P., & Pomerol, J. C. (1999). Contextual Knowledge Sharing and Cooperation in Intelligent Assistant Systems. *Le travail humain*, 223-246.
- Brown, P. J., Bovey, J. D., & Chen, X. (1997). Context-Aware Applications: from The Laboratory to the Marketplace. *IEEE Personal Communications*, 4(5), 58-64.
- Buyurgan, N., & Lehlou, N. (2015). A Terrain Risk Assessment Method for Military Surveillance Applications for Mobile Assets. *Computers & Industrial Engineering*, 88, 88-99.
- Canon-Bowers, J., Salas, E., & Converse, S. (1993). Shared Mental Models in Expert Team.
- Chen, A. (2005, May). Context-aware Collaborative Filtering System: Predicting the User's Preference in the Ubiquitous Computing Environment. In *International Symposium on Location-and Context-Awareness* (pp. 244-253). Springer, Berlin, Heidelberg.
- Clark, H. H. (1996). *Using language*. Cambridge University Press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in Communication. In L. B. Resnick, J. M.

- Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). American Psychological Association. <https://doi.org/10.1037/10096-006>
- Clark, H. H. & Marshall, C. R. (1981). Definite Knowledge and Mutual Knowledge. In Aravind K. Joshi, Bonnie L. Webber & Ivan A. Sag (eds.), *Elements of Discourse Understanding*. Cambridge, UK: Cambridge University Press. pp. 10–63.
- Clark, H. H., & Schaefer, E. F. (1987). Collaborating on Contributions to conversations. *Language and Cognitive Processes*, 2(1), 19-41.
- Davis, A. B., Moore, M. M., & Storey, V. C. (2003, June). Context-Aware Communication for Severely Disabled Users. In *Proceedings of the 2003 Conference on Universal Usability* (pp. 106-111).
- Demir, M., McNeese, N. J., Cooke, N. J., Ball, J. T., Myers, C., & Freiman, M. (2015, September). Synthetic Teammate Communication and Coordination with Humans. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 59, No. 1, pp. 951-955). Sage CA: Los Angeles, CA: SAGE Publications.
- Dey, A. K. (2001). Understanding and using Context. *Personal and Ubiquitous Computing*, 5(1), 4-7.
- Dey, A. K., & Abowd, G. D. (1999). *Towards a Better Understanding of Context and Context-Awareness*. Georgia Institute of Technology.
- Dey, A. K., & Abowd, G. D. (2000, September). Cybreminder: A Context-Aware System for Supporting Reminders. In *International Symposium on Handheld and Ubiquitous Computing* (pp. 172-186). Springer, Berlin, Heidelberg.
- Dey, A. K., Abowd, G. D., & Salber, D. (2001). A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of context-Aware Applications. *Human-Computer Interaction*, 16(2-4), 97-166.
- Dourish, P. (2004). What we Talk about when we Talk about Context. *Personal and Ubiquitous Computing*, 8(1), 19-30.
- Endsley, M. R. (1995). Measurement of Situation Awareness in Dynamic Systems. *Human factors*, 37(1), 65-84.
- Estival, D., & Molesworth, B. (2012). Radio Miscommunication: EL2 Pilots in the Australian General Aviation Environment. *Linguistics and the Human Sciences*, 5(3), 351-379.
- Ferreira, F., & Patson, N. D. (2007). The 'Good Enough' Approach to Language Comprehension. *Language and Linguistics Compass*, 1(1-2), 71-83.
- Fischer, G. (2012, May). Context-Aware Systems: The 'Right' Information, at the 'right' Time, in the 'right' Place, in the 'right' Way, to the 'Right' Person. In *Proceedings of the International Working Conference On Advanced Visual Interfaces* (Pp. 287-294).
- Fletcher, C. R., & Chrysler, S. T. (1990). Surface Forms, Textbases, And Situation Models: Recognition Memory For Three Types Of Textual Information. *Discourse processes*, 13(2), 175-190.
- Fogarty, J., Lai, J., & Christensen, J. (2004). Presence versus Availability: The Design and

- Evaluation of a Context-Aware Communication Client. *International Journal of Human-Computer Studies*, 61(3), 299-317.
- Franklin, D., & Flaschbart, J. (1998, March). All Gadget and no Representation Makes Jack a Dull Environment. In *Proceedings of the AAAI 1998 Spring Symposium on Intelligent Environments* (pp. 155-160).
- Fraundorf, S. H., Benjamin, A. S., & Watson, D. G. (2013). What Happened (and what did not): Discourse Constraints On Encoding Of Plausible Alternatives. *Journal Of Memory And Language*, 69(3), 196-227.
- Fraundorf, S. H., Watson, D. G., & Benjamin, A. S. (2010). Recognition memory reveals just how CONTRASTIVE contrastive accenting really is. *Journal of memory and language*, 63(3), 367-386.
- Freiman, M., Caisse, M., Ball, J., Halverson, T., & Myers, C. (2018, June). Empirically identified gaps in a situation awareness model for human-machine coordination. In *2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)* (pp. 110-116). IEEE.
- Garnham, A., Oakhill, J., & Johnson-Laird, P. N. (1982). Referential continuity and the coherence of discourse. *Cognition*, 11(1), 29-46.
- Glenberg, A. M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. *Journal of memory and language*, 26(1), 69-83.
- Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41-58). Brill.
- Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of verbal learning and verbal behavior*, 13(5), 512-521.
- Hempelmann, C. F., Dufty, D., McCarthy, P. M., Graesser, A. C., Cai, Z., & McNamara, D. S. (2005). Using LSA to automatically identify givenness and newness of noun phrases in written discourse. In *Proceedings of the 27th annual conference of the Cognitive Science Society* (pp. 941-946).
- Henricksen, K. (2003). A framework for context-aware pervasive computing applications.
- Hull, R., Neaves, P., & Bedford-Roberts, J. (1997, October). Towards situated computing. In *Digest of papers. first international symposium on wearable computers* (pp. 146-153). IEEE.
- Iwatani, Y. (1998, June 11). *Love: Japanese style*. Wired. Retrieved October 10, 2022, from <https://www.wired.com/1998/06/love-japanese-style/>
- Jang, S., & Woo, W. (2003, June). Ubi-UCAM: A unified context-aware application model. In *International and Interdisciplinary Conference on Modeling and Using Context* (pp. 178-189). Springer, Berlin, Heidelberg.
- Jiang, X., Chen, N. Y., Hong, J. I., Wang, K., Takayama, L., & Landay, J. A. (2004, April). Siren: Context-aware computing for firefighting. In *International Conference on Pervasive Computing* (pp. 87-105). Springer, Berlin, Heidelberg.
- Kane, S. K., Linam-Church, B., Althoff, K., & McCall, D. (2012, October). What we talk about:

- designing a context-aware communication tool for people with aphasia. In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility* (pp. 49-56).
- Kayes, A. S. M., Han, J., & Colman, A. (2015). An ontological framework for situation-aware access control of software services. *Information Systems*, 53, 253-277.
- Kintsch, W. (1998). The representation of knowledge in minds and machines. *International Journal of Psychology*, 33(6), 411-420.
- Kintsch, W., Welsch, D., Schmalhofer, F., & Zimny, S. (1990). Sentence memory: A theoretical analysis. *Journal of Memory and language*, 29(2), 133-159.
- Kleinman, D. L., & Serfaty, D. (1989, April). Team performance assessment in distributed decision making. In *Proceedings of the symposium on interactive networked simulation for training* (pp. 22-27). Orlando, FL: University of Central Florida.
- Kokinov, B. (1999, September). Dynamics and automaticity of context: A cognitive modeling approach. In *International and interdisciplinary conference on modeling and using context* (pp. 200-213). Springer, Berlin, Heidelberg.
- Kopp, S., & Krämer, N. (2021). Revisiting human-agent communication: The importance of joint co-construction and understanding mental states. *Frontiers in Psychology*, 12, 580955.
- Lamothe, D. (2014, September 4). Investigation: Friendly fire airstrike that killed U.S. Special Forces was avoidable. *Washington Post*.
- Marlow, S. L., Lacerenza, C. N., Paoletti, J., Burke, C. S., & Salas, E. (2018). Does team communication represent a one-size-fits-all approach?: A meta-analysis of team communication and performance. *Organizational behavior and human decision processes*, 144, 145-170.
- Marmasse, N., & Schmandt, C. (2000, September). Location-aware information delivery with commotion. In *International symposium on handheld and ubiquitous computing* (pp. 157-171). Springer, Berlin, Heidelberg.
- Mata, A. (2020). An easy fix for reasoning errors: Attention capturers improve reasoning performance. *Quarterly Journal of Experimental Psychology*, 73(10), 1695-1702.
- Matsumoto, S., & Riek, L. D. (2022). Shared Control in Human Robot Teaming: Toward Context-Aware Communication. *arXiv preprint arXiv:2203.10218*.
- Menenti, L., Pickering, M. J., & Garrod, S. C. (2012). Toward a neural basis of interactive alignment in conversation. *Frontiers in human neuroscience*, 6, 185.
- McChrystal, G. S., Collins, T., Silverman, D., & Fussell, C. (2015). *Team of teams: New rules of engagement for a complex world*. Penguin.
- McCrudden, M. T., & Schraw, G. (2007). Relevance and goal-focusing in text processing. *Educational psychology review*, 19(2), 113-139.
- McNeese, N. J., Demir, M., Cooke, N. J., & Myers, C. (2018). Teaming with a synthetic teammate: Insights into human-autonomy teaming. *Human factors*, 60(2), 262-273.

- Miluzzo, E., Lane, N. D., Fodor, K., Peterson, R., Lu, H., Musolesi, M., ... & Campbell, A. T. (2008, November). Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM conference on Embedded network sensor systems* (pp. 337-350).
- Mustajoki, A. (2012). A speaker-oriented multidimensional approach to risks and causes of miscommunication. *Language and dialogue*, 2(2), 216-243.
- Myers, C., Ball, J., Cooke, N., Freiman, M., Caisse, M., Rodgers, S., ... & McNeese, N. (2018). Autonomous intelligent agents for team training. *IEEE Intelligent Systems*, 34(2), 3-14.
- Mynatt, E. D., Back, M., Want, R., Baer, M., & Ellis, J. B. (1998, January). Designing audio aura. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 566-573).
- Norberg, Kole A. "Avoiding Miscomprehension: A Metacognitive Perspective for how Readers Identify and Overcome Comprehension Failure." PhD diss., University of Pittsburgh, 2022.
- Norberg, K. A., & Fraundorf, S. H. (2021). Memory benefits from contrastive focus truly require focus: evidence from clefts and connectives. *Language, Cognition and Neuroscience*, 36(8), 1010-1037.
- Paleja, R. (2022). Mutual Understanding in Human-Machine Teaming. (Pre-Print)
- Pascoe, J. (1997, January). The stick-e note architecture: extending the interface beyond the user. In *Proceedings of the 2nd international conference on Intelligent user interfaces* (pp. 261-264).
- Radvansky, G. A., Gerard, L. D., Zacks, R. T., & Hasher, L. (1990). Younger and older adults' use of mental models as representations for text materials. *Psychology and Aging*, 5(2), 209.
- Rakas, J., & Yang, S. (2007). Analysis of multiple open message transactions and controller-pilot miscommunications. In *7th USA/Europe Air Traffic Management R&D Seminar, Barcelona*.
- Ranganathan, A., & Lei, H. (2003). Context-aware communication. *Computer*, 36(4), 90-92.
- Rothwell, C., Gluck, K., Ayres, D., Mullins, B., Mobley, F., Harmer, T., & Romigh, G. (2022, September). *Red Flag – Rescue communications analysis dashboard*. Presentation for 414th Combat Training Squadron, Detachment 1. Davis-Monthan Air Force Base, AZ.
- Ryan, N. S., Pascoe, J., & Morse, D. R. (1998). Enhanced reality fieldwork: the context-aware archaeological assistant. In *Computer applications in archaeology*. Tempus Reparatum.
- Savitsky, K., Keysar, B., Epley, N., Carter, T., & Swanson, A. (2011). The closeness-communication bias: Increased egocentrism among friends versus strangers. *Journal of experimental social psychology*, 47(1), 269-273.
- Schilit, B. N., Hilbert, D. M., & Trevor, J. (2002). Context-aware communication. *IEEE Wireless Communications*, 9(5), 46-54.
- Schilit, B. N., & Theimer, M. M. (1994). Disseminating active map information to mobile

- hosts. *IEEE network*, 8(5), 22-32.
- Schmalhofer, F., & Glavanov, D. (1986). Three components of understanding a programmer's manual: Verbatim, propositional, and situational representations. *Journal of memory and language*, 25(3), 279-294.
- Schmidt, A. (2003). *Ubiquitous computing-computing in context*. Lancaster University (United Kingdom).
- Schwarzschild, R. (1999). GIVENness, AvoidF and other constraints on the placement of accent. *Natural language semantics*, 7(2), 141-177.
- Shishkov, B., Larsen, J. B., Warnier, M., & Janssen, M. (2018, July). Three categories of context-aware systems. In *International Symposium on Business Modeling and Software Design* (pp. 185-202). Springer, Cham.
- Siewiorek, D. P., Smailagic, A., Furukawa, J., Krause, A., Moraveji, N., Reiger, K., ... & Wong, F. L. (2003, October). SenSay: A Context-Aware Mobile Phone. In *ISWC* (Vol. 3, p. 248).
- Skaltsas, G., Rakas, J., & Karlaftis, M. G. (2013). An analysis of air traffic controller-pilot miscommunication in the NextGen environment. *Journal of Air Transport Management*, 27, 46-51.
- Tajima, A. (2004). Fatal miscommunication: English in aviation safety. *World Englishes*, 23(3), 451-470.
- Urban, J. M., Bowers, C. A., Monday, S. D., & Morgan Jr, B. B. (1995). Workload, team structure, and communication in team performance. *Military Psychology*, 7(2), 123-139.
- Van den Broek, P., Risdén, K., & Husebye-Hartmann, E. (1995). The role of readers' standards for coherence in the generation of inferences during reading. In *Part of this research was reported at the Annual Meeting of the American Educational Research Assn, Chicago, 1991*. Lawrence Erlbaum Associates, Inc.
- van Kasteren, A., & Vredenburg, M. (2022, July). Personalized, context-aware communication in multimodal public transport. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization* (pp. 326-330).
- Vieira, V., Tedesco, P., & Salgado, A. C. (2011). Designing context-sensitive systems: An integrated approach. *Expert Systems with Applications*, 38(2), 1119-1138.
- Volpe, C. E., Cannon-Bowers, J. A., Salas, E., & Spector, P. E. (1996). The impact of cross-training on team functioning: An empirical investigation. *Human factors*, 38(1), 87-100.
- Whitaker, B. (2017, November 15). Why were 5 U.S. soldiers killed by an American bomber in Afghanistan? *CBS News*. Retrieved October 6, 2022, from <https://www.cbsnews.com/news/why-were-five-u-s-soldiers-killed-by-an-american-bomber-in-afghanistan/>
- Winograd, T. (2001). Architectures for context. *Human-Computer Interaction*, 16(2-4), 401-419.
- Yang, C. L., Perfetti, C. A., & Schmalhofer, F. (2007). Event-related potential indicators of text integration across sentence boundaries. *Journal of Experimental Psychology: Learning*,

- Memory, and Cognition*, 33(1), 55.
- Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological bulletin*, 127(1), 3.
- Zimmermann, A., Lorenz, A., & Oppermann, R. (2007, August). An operational definition of context. In *International and interdisciplinary conference on modeling and using context* (pp. 558-571). Springer, Berlin, Heidelberg.
- Zimmermann, M. (2008). Contrastive focus and emphasis. *Acta Linguistica Hungarica*, 55(3-4), 347-360.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological bulletin*, 123(2), 162.

7.0 LIST OF ACRONYMS, ABBREVIATIONS AND SYMBOLS

ABC	Actively Building Context
NLP	Natural Language Processing
HMT	Human-Machine Teaming
JTAC	Joint Terminal Attack Controller
GPC	Global Positioning System
SIREN	Secure Integrated Response Electronic Notification
AFRL	Air Force Research Laboratory
ACT-R	Adaptive Control of Thought-Rational
UAV	Unmanned Aerial Vehicle