

Team Context-Aware Collaborative AI Assistants

Micael Vignati, Matthew Johnson

The Florida Institute for Human and Machine Cognition
mvignati@ihmc.org, mjohnson@ihmc.org

Abstract

True collaboration or teamwork involves being aware of the team and its goals, adapting to the changing situations, communicating relevant information, identifying breakdowns in common ground and repairing them, all to improve joint outcomes. At the heart of this type of collaborative agent are the capabilities needed to understand context and determine relevance. We introduce four key capabilities that enable our collaborative agents to be effective collaborators by understanding context and determining relevance.

Introduction

The concept of collaborative AI assistants has been around for quite some time (Bayus, Jain, and Rao 1997; Clancey 2004; Ferguson, Allen et al. 1998; Hexmoor and Vaughn 2002). Despite the advances of new technology, such as Siri and large language models, the dream of effective personal AI assistants that help us coordinate both our work and our lives remains elusive. Many efforts strive to make new systems “smarter” by answering questions correctly, by doing their work more reliably, or requiring less assistance. These efforts will make the agents more capable, but not more collaborative.

True collaboration or teamwork involves being aware of the team and its goals, adapting to the changing situations, communicating relevant information, identifying breakdowns in common ground and repairing them—all to improve joint outcomes. At the heart of this type of collaborative agent is the capability to understand context and determine relevance.

Many domains of interest today are faced with too much data, yet not enough information. Data is inherently meaningless without context. Context is necessary to determine relevance, interpret data effectively, and thus it is key to framing and understanding. As such, context underpins all decision making. For an AI agent to be an effective collaborator or teammate, it needs to understand context.

Relevance is information that will positively impact the activity (Sperber and Wilson 1986). Understanding contextual relevance is complicated by the natural dynamics of

many domains which are continuously changing. It is further complicated by distributed and asynchronous teams as each team member has different awareness, different roles, and different decision-making requirements. Any collaborative AI assistant will need to understand individual context as well as the context of any other team members to effectively determine relevance and facilitate collaboration.

As part of an ONR effort, we have been developing collaborative AI-enabled (CAI) agents to support data-intensive mission planning and execution for disaster response. Our goal is to develop CAI agents with the capability to understand context and determine relevance to help assess, curate, process, and manage data in support of decision making. We will first discuss context and relevance and then introduce the four key capabilities we are building to enable our CAI agents to be effective collaborators.

Context

Context is commonly thought of as “any information that can be used to characterize the situation of an entity” (Dey 2001). However, this interpretation could lead to all data being considered context, which is impractical for real systems. We propose that there are two key elements of context that are practical to track and model: *activity context* and *team context*.

Activity context refers to the set of goals and the plans to achieve those goals. Important components of activity context include the structure of the activity, which is shaped by several factors: information flow, resource constraints, and state requirements (pre-conditions). Additionally, the type of activity—such as additive, conjunctive, or disjunctive (Steiner 1972)—provides meaningful context. All of these factors introduce interdependencies that are crucial for interpreting data. Finally, the state of progress within an activity serves as another important contextual element.

Team context pertains to the composition and status of team members. Composition identifies who is part of the group. Availability and capabilities determine who can participate or contribute to an activity. Roles define responsibilities and expected behaviors. Participation refers to who actually engages in the activity and their current status. A more in-depth discussion of these aspects can be found in the description of the 4S Interdependence Framework (Johnson and Vera 2019).

Relevance

Relevance refers to information that can positively impact an activity (Sperber and Wilson 1986). A positive impact may involve an increase in value, such as making the activity faster or more accurate, or reducing cost by preventing errors or conserving resources. Conversely, information about potential negative impacts is also relevant, as it enables proactive measures that prevent harm, thereby producing a positive effect through the avoidance of negative outcomes. Information that neither directly nor indirectly contributes to improving the activity is considered irrelevant.

Collaborative Capabilities

Our goal is to develop CAI agents with the capability to understand context and determine relevance to help assess, curate, process, and manage data in support of decision making. We propose four collaborative capabilities to enable this:

- Real-time Joint Activity Context Tracking
- Context-Aware Communication Engine
- Personalized Knowledge Inference & Information Filtering
- Personalized CAI Networked Helper

Real-time Joint Activity Context Tracking

We leverage a dynamic Joint Activity Graph (JAG) (Vignati et al. 2022) to enable real-time tracking of joint activity. JAGs can be used to capture both the activity context and the team context. A JAG is a hierarchical decomposition of a given joint activity. It defines the goals and the activity search space for the agents. It also captures the information flow requirements within the activity. Together, this captures the structural constraints of the activity, both static and dynamic.

JAGs can be used to track activity status at runtime as well as team participation, providing team context. JAGs integrate real-time data feeds and event-based triggers to update team context when appropriate. This provides continuous, but deliberate, updates and reflects the team’s goals and each team member’s current tasks and status. It enables distributed teams to gain immediate situational awareness, reducing miscommunication and enabling rapid, coordinated responses during operations.

Information is data in context, and JAGs provide structured context in support of collaboration. The crucial aspect is understanding the context of a given instance of activity. When seeking information, we do not want a general answer, but an answer customized and tailored to the current context of this specific team and mission.

Context-Aware Communication Engine

We propose that it is one’s understanding of their current context that is used to determine relevance, as depicted in Figure 1.

We utilize JAG-based context models that include role-based filtering algorithms to prioritize and route messages based on live operational data and individual user profiles

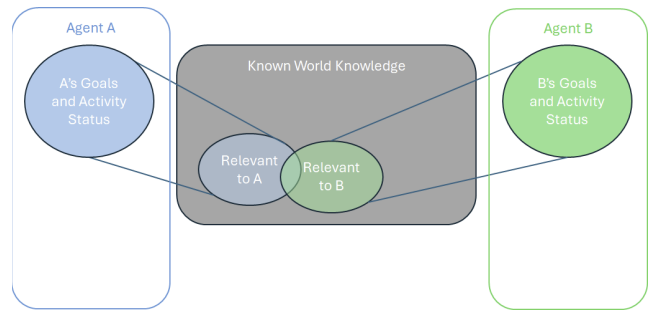


Figure 1: An agent’s context (goals, activity status and team status) provides a lens for interpreting available knowledge to determine what is relevant.

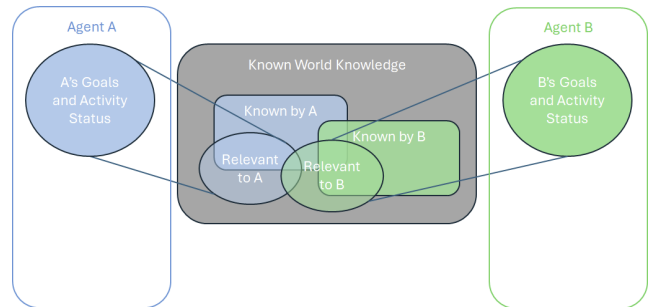


Figure 2: We track user’s knowledge to further refine the scope of what is relevant to communicate.

(roles and preferences). This automatically shares targeted and relevant information based on the current operational context and the specific role of each team member. The approach ensures that critical information reaches the right team members as soon as possible, improving team decision making speed, reducing the risk of errors, and improving efficiency.

Personalized Knowledge Inference & Information Filtering

Understanding what a teammate knows is non-trivial. In our domain, we leverage the fact that participants use a web-based interface to coordinate, which provides a means to track and estimate, within reason, what a given person is aware of.

We utilize the JAG-based user context model to assess and selectively filter information from external knowledge sources, tailoring information to each user’s current knowledge state and activity status, as depicted in Figure 2.

This is accomplished by the CAI assistant tracking what each user knows and selectively filtering incoming information to highlight only the most relevant and actionable insights. This approach reduces information overload by filtering out non-essential data, ensuring that every team member receives concise, actionable, and timely mission-critical updates.

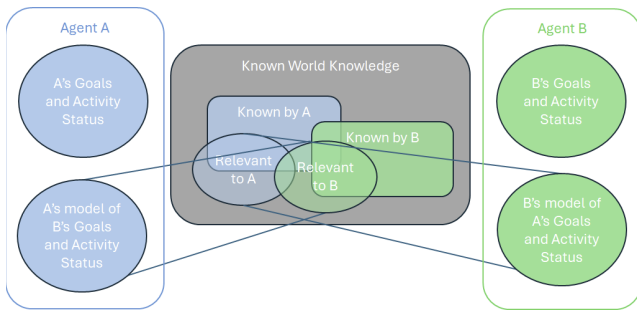


Figure 3: Modelling knowledge of other team members to support inter-team collaboration and coordination

Personalized CAI Networked Helper

Thus far, we have modeled individual context and knowledge to assist that specific individual in their decision-making. We can leverage the exact same approach to support collaboration across the team. Our CAI models not only the person's context and awareness, but also what that person understands about their teammates (i.e., Theory of Mind), as depicted in Figure 3.

By modelling each team member, they can look for misalignments in common ground to help avoid miscoordination. Once identified, our assistants can facilitate resolutions and ensure teams manage their interdependence efficiently. This targeted approach enhances communication clarity, reduces overload, and enables fast, data-driven decisions in the field.

By doing so, we enable our CAI assistant to function as an intelligent personal assistant that actively connects to the broader communication network, aggregating data from multiple sources, including other team members, to deliver tailored, role-specific alerts and insights. It integrates and provides multi-source data aggregation and modeling with machine learning-driven personalization, delivering contextual alerts through a web client interface.

This approach empowers each team member with a dedicated, networked helper that keeps them updated and in sync with critical information, even in high-pressure, distributed operational environments. This enables our CAI agents to facilitate effective teamwork and coordination among team members.

Discussion

These four capabilities provide our CAI agents with the capability to understand context and determine relevance. An important consideration of this design is that all aspects of this (e.g., the relevance of information and team member knowledge) are dynamic and will change over time.

Good teams reliably produce good mission outcomes by adapting to their changing situations. Individuals on the team know how to adapt through awareness of relevant information (about the team, the activity, and the environment) and application of that information with an understanding of activity and team context (i.e., team context-aware decision-making). The team must communicate to share relevant con-

textual understanding and ensure all team members have sufficient information for a coordinated and coherent effort while managing the cost of collaboration.

Development of truly collaborative AI-enabled assistants requires such agents to track, understand, and utilize team context to facilitate individual and team performance in collaborative activities, and this is what our design provides.

Conclusion

Information is common to all activities. Decision making is based on information and its value; however, the value of information can only be determined in context. Therefore, context is fundamental to effective decision-making. We need better ways to capture and represent context. It is essential for humans to make good decisions, and it is equally important to CAI agents trying to inform and assist those people.

Collaborative AI-enabled assistants should provide information, answers, or suggestions customized and tailored to the current context of this specific team and mission.

Acknowledgments

This material is based upon work sponsored by the Office of Naval Research under grant Agreement number N00014-23-1-2886.

References

- Bayus, B. L.; Jain, S.; and Rao, A. G. 1997. Too Little, Too Early: Introduction Timing and New Product Performance in the Personal Digital Assistant Industry. *Journal of Marketing Research*, 34(1): 50–63.
- Clancey, W. J. 2004. Roles for Agent Assistants in Field Science: Understanding Personal Projects and Collaboration. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2): 125–137.
- Dey, A. K. 2001. Understanding and Using Context. *Personal and Ubiquitous Computing*.
- Ferguson, G.; Allen, J. F.; et al. 1998. TRIPS: An Integrated Intelligent Problem-Solving Assistant. In *AAAI/IAAI*, 567–572.
- Hexmoor, H.; and Vaughn, J. T. 2002. Computational Adjustable Autonomy for NASA Personal Satellite Assistants. In *Proceedings of the ACM Symposium on Applied Computing*, 21–26. Association for Computing Machinery (ACM).
- Johnson, M.; and Vera, A. H. 2019. No AI Is an Island: The Case for Teaming Intelligence. *AI Magazine*, Spring: 16–28.
- Sperber, D.; and Wilson, D. 1986. *Relevance: Communication and Cognition*. Harvard University Press.
- Steiner, I. D. 1972. *Group Process and Productivity*. Academic Press.
- Vignati, M.; Johnson, M.; Bunch, L.; Carff, J.; and Duran, D. 2022. Interdependence Design Principles in Practice. *Frontiers in Physics*, 10(Interdisciplinary Approaches to the Structure and Performance of Interdependent Autonomous Human Machine Teams and Systems (A-HMT-S)).