Using Context and Robot-Human Communication to Resolve Unexpected Situational Conflicts

Wayne W. Zachary
CHI Systems, Inc.
Plymouth Meeting, PA, USA
wzachary@chisystems.com

Taylor J. Carpenter
CHI Systems, Inc.
Plymouth Meeting, PA, USA
tcarpenter@chisystems.com

Abstract— **While efforts to develop cognitive abilities for robots have made progress from the perspective of goal-directed task performance, research has shown that additional cognitive capabilities are needed to enable robots to interact, cooperate, and act as teammates with humans. In particular, robots need additional teamwork and coordination knowledge and an ability to apply this knowledge to a model of context that is at least homologous to the context models that people use in reasoning about environmental interactions. The Context-Augmented Robotic Interface Layer (CARIL) provides a robot with a cognitively-motivated computational capability for situation assessment and situational adaptation. CARIL is used to analyze and develop context-based reasoning strategies that allow a robot to coordinate its behavior and spatial movements with humans when they are working on shared tasks and/or in shared space. Both communication-free and communications approaches are addressed and tested in a simulated environment.**

Keywords—Context, situational awareness, action compliance, cognitively-inspired technology,

I. INTRODUCTION

A substantial body of research exists on developing cognitive capabilities for robots and other autonomous devices. The overwhelming majority of that work has focused on enabling purposive cognition in the form of task- and goal-oriented problem solving (e.g., [1],[2]). Such research has sought to make a robot competent in performing individual work tasks autonomously, e.g., as a surrogate for human workers. The resulting cognitive systems have met with measured, if modest success [3], but have not translated well when applied to the broader desire to use these entities as teammates and collaborators in applications such as human robotic interaction (e.g. [4-10]). Research and application efforts in these areas have found that the knowledge and cognitive capabilities needed to perform a task individually (i.e., purposive task cognition) are not sufficient to allow a solitary task-performer to carry out that same task cooperatively or collaboratively [11], [12]. In particular, two additional cognitive capabilities have emerged as necessary for teamwork and cooperation:

- Understanding context: having a shared representation of the cooperative task and setting – social, physical, and environmental -- in which that task is being performed; and
- Teamwork/coordination knowledge: having knowledge, both general and domain-dependent, about how to coordinate activity in completing the task and applying it to the local context.

This research was supported by Contract NNX14CJ38P from the NASA Johnson Space Center and in part by contract N00014-15-C-5068 from the Office of Naval Research. We gratefully acknowledge their support. This paper reflects only the ideas and conclusions of the authors and not the sponsoring agencies; any errors or omissions are ours alone.

The (as yet, unmet) modeling challenge has been to translate these high-level concepts into specific modeling methods and computational/simulation tools that can be used to design and build robots and/or interactive software agents with teamwork and cooperative capabilities.

We believe that a solution to this challenge lies in using the cognitive science of human-human collaboration to re-think the human-robot interaction problem. There are two aspects to this approach. The first is understanding that context is a key emergent feature of the common underlying cognitive architecture that all people share. As Suchman [13] noted in her seminal studies of human-machine interaction, human cognition is situated in the social and physical context, and we (people) use our cognitive representation of the situation in acting and *inter*-acting (including the pragmatic use of language). This common understanding simplifies the communication that is needed for effective interaction, and establishes (usually implicitly) frames for expectations and predictions of each other's actions and utterances. The second aspect of this approach is understanding that robots do *not* share this common architecture, and are unable to implicitly participate in interactions that presume it. As Hoffman and Woods [14] put it, "*machines do not know that they are in the world that they have within themselves as a model of that world*". Thus, if robots are to be given any chance of being intelligent subordinates (much less collaborators or cooperative teammates), then they need to have access to some representation of the situational context *as the humans involved would understand it.*

The authors and colleagues have been pursuing a line of research and technology development to provide robots with such a capability. The goal of this approach is not to precisely model human situational awareness and reasoning but rather to create a computational representation that is compatible (enough) that the machine agents can cooperate and coordinate in a human-like, and more importantly human-understandable, way. Zachary et al. [15] introduces the general theory underlying this approach, and [16] defines a formal computational system to implement and execute it in software. Here these prior results are applied to problems of robot-human coordination where the robot(s) and human(s) share tasks and/or physically-constrained workspaces, such as warehouses, factory floors, construction sites and engineering facilities. The particular example used here is that of robots and human astronauts working together in a space-habitat, though the approach is general and could be applied to any case from the other domains listed above.

II. CARIL

The Context-Augmented Robotic Interaction Layer (CARIL) is a generalized computational architecture that gives a robot, in a mixed human-robot workplace, a human-like representation understanding of its work situation and environment, and the ability to reason about context in order to adapt its behavior, to that of the humans around it. We have called this capability "*action compliance*" [17].

A. Overview of CARIL

CARIL consists of five interconnected components:

- A symbolic declarative representation of the (dynamic) situation in which the robot is working;
- A set of self-activating and self-organizing procedural knowledge elements that continuously work to build and maintain the situational representation in the light of external information, and internal changes to the situation representation;
- A set of self-activating and self-organizing adaptive reasoning knowledge elements that recognize current and (ideally) future situational evolutions in the declarative situational representation that could require the robot to develop an action compliant response to unexpected human activities, and generate the needed action-compliant behaviors;
- A behavioral action director, that can translate intended robot actions into directives that are executed by lower-level physical controllers; and
- A perceptual channel, that accepts processed inputs from physical sensors and/or environmental data streams.

The declarative situational representation is the central feature because it is structured to resemble the semantic organization of human situation awareness. Its structure is derived from the body of work on human situation awareness and situational understanding, and particularly from the cross-domain structure identified in the work of Mica Endsley and colleagues (e.g. [18], [19]). The situational representation contains four levels of abstraction and function:

- *Perception*, in which the existence, status, and/or attributes, of relevant physical or informational in the environment are perceived as distinct entities following the sensation of information about them;
- *Significance*, in which more abstract and/or ontological characteristics and relationships among or of perceived entities are identified or constructed through a reasoning process. These elements of comprehension can include information on how the perceived elements can impact the robot's situational goals;
- *Expectations*, in which possible or expected future actions of perceived elements or significance-level abstractions of them are projected forward in time; and,
- *Plans, conflicts and adaptations,* in which potential conflicts between the robot's plans and the actions of human co-workers are identified, reasoned about, and used to develop adaptations of the robot's plans and actions to comply with human behaviors.

Here, we focus specifically on work situation in which there is general plan for both human and robot activity, but where actual (human) actions can deviate unexpectedly from the general plan (for whatever reason).

As specific information items are perceived by CARIL, they are internalized into the situational representation. This perceptual (internalization) process may stimulate procedural knowledge elements to self-activate and reason about the new information, e.g., to create relationships between new and existing information, to discern aspects of significance and post them, to the significance level, etc. Those internally driven changes can lead to other procedural knowledge elements becoming activated and making further representational changes, e.g., to make projections about future behavior on the expectations level. In this way, a dynamic representation of the current situation is built and maintained over time. In computational terms, this overall process is an example of the broad class of Pandemonium architectures, first suggested by Selfridge [20] and widely used throughout computer science and artificial intelligence.

To these classical situation awareness levels, CARIL adds a fourth level, that of conflicts and adaptations. This is a level at which the third CARIL component – the adaptive reasoning elements – operates, identifying potential future situations (i.e., conflicts) that could require action compliance. In CARIL, the reasoning for both situation awareness maintenance and action compliance uses additional sets of background knowledge. These are knowledge about the physical layout of the workspace, and knowledge about the general work plans or schedules of the humans and the robot in the workplace.

It is important to note that CARIL is a *cognitive* system for a robot, and its not intended to solve problems that involve sensation and perception (e.g., identifying objects in the environment), or action implementation (e.g., translating intended actions, such as "move to x,y and pick up the laptop computer). These are separate areas of research and largely involve hardware-in-loop technology. However, we assert that problems of dynamic, context-sensitive action compliance would remain and would require technology like CARIL even if the perception and action implementation problems were fully solved.

B. Using Situational Awareness to Direct Robot Work and Action-compliance

CARIL approaches action compliance as an aspect of robot work behavior. That is, it presumes the robot is there (in our example, in a space habitat) to carry out a work plan of is own, but must also do so in a way that does not interfere with human astronaut activities. Thus, CARIL directs the robot(s) it is controlling to complete its (their) daily work plan, but is prepared to deviate from that plan to avoid interfering with the work or movements of the astronauts, even when that human work or movements deviate from the posted work schedule.

Put differently, robot actions in CARIL all stem from one of two purposes:

- Plan compliance (to comply with the work plan for the robot), and/or
- Action compliance (to adapt its activities to human actions that are unexpected or that deviate from the general plan in time, space, or detail.

The behaviors that CARIL generates, as solutions to a potential conflict, can be either physical movements of the robot, dialogs with one or more humans, or both. CARIL can choose to initiate and engage in a dialog with a person primarily to confirm or clarify assumptions that underlie a presumed cause of a conflict, or to resolve or solicit a specific solution to a conflict (see below). While CARIL could engage in communication any time there is an inferred conflict, there are reasons why it might not want to. These are primarily for the benefit of the people involved – not to disturb a busy person at a critical moment, or to avoid communicating about a subject that the human might consider obvious and thus annoying.

Thus, these communication behaviors are treated in CARIL as what linguistics calls "speech acts", communications made to serve a specific pragmatic function (see [21], [22], regarding the concept of speech acts, and [23] for an overview of language pragmatics). The point with this approach is that the communicative acts are initiated and made context-specific as part of the context reasoning process that manages the overall non-interference action compliance.

III. An Example Case

The above approach to action compliance is being operationalized in a space exploration use case, in which four human astronauts are working in a space habitat (analogous to the international space station) with the support of an anthropomorphic robot. The astronauts have a general schedule in which they move around the habitat doing routine activities (eating, sleeping, exercising, attending meetings) and work tasks that include both experiments and repair/maintenance tasks. The robot also has a daily schedule of activities that involve independent work tasks (typically maintenance) that are done away from the astronauts, to avoid any possibility of accidental collisions between the robot and astronauts. The action compliance during these tasks involves identifying situations where an astronaut may be deviating in time and/or space from her/his general activity plan.

In CARIL, this kind of action-plan compliance reasoning is an on-going reasoning cycle in which an (internal) context representation is built using real-time tracking data feeds on all robots and astronauts in the environment combined with the current daily plans. Several types of deviation from the plan are identified and categorized by CARIL as part of this reasoning process:

- Astronaut is staying late at a planned task (which then turns into astronaut is leaving late for a next task);
- Astronaut is leaving early from planned task; and
- Astronaut is leaving a planned task in a way that can't be readily diagnosed as one of the above.

Detecting these situations is of interest to non-interference because each has the potential to lead directly to a future situation in which the robot's compliance to its own plan can lead it to interfere with an astronaut.

It should be noted that in cases 1-3, CARIL can not actually know why the astronaut is at the current location, only that it its not where the astronaut is scheduled to be. In each of these cases, CARIL makes a plausible assumption, termed a Working Hypothesis or WH, which is the intention that is listed above for categories 1) through 3). In category 4, CARIL makes no assumption about the reason for the astronaut's departure from the schedule. The reasoning that follows differs in cases 1) through 3), collectively termed presumed intention cases, versus case 4), termed an unknown intention case. However, in all cases, the reasoning that follows is based on a two stage process of first, reasoning whether an interference situation would follow if the robot did not deviate from its plan, and second (if an interference were projected), reasoning about what action(s) could be taken to avoid the interference.

For example, in case 1), CARIL uses the presumed intent to reason forward in time to determine if the astronaut's continued presence in the current location would interfere with a future movement of the robot. If so, CARIL then reasons about how to avoid it. Without communicating, the robot defaults to a stance of staying in its current location even after its current work assignment is done, and wait until the astronaut moves to the next planned task. In cases 2) and 3), the astronaut is moving so CARIL uses the presumed intent (moving to the next work location on the astronaut's plan) as a presumed movement path, and then analyzes it for conflicts with the robot's planned (stationary) location or planned movements to develop an action compliance strategy. In the non-communication case, the strategies each involve applying heuristics to make assumptions (termed Working Hypotheses) about the intent of the astronaut, and using them to plan and execute adaptive action-compliant actions. If and when those heuristics and assumptions prove correct, the actions will result in non-interference. Otherwise, though, the actions could lead to a case of interference, despite the situationally-based reasoning.

When communication with astronauts is not permitted, CARIL simply identifies patterns of information (i.e., situations) in the context representation which are instances of any of the above four types of the type of situation. It then creates a concept of a potential conflict on the plans, adaptations, and conflicts panel of the context representation. Once such a potential concept is created, it further reasons to see if an actual conflict is likely (e.g., astronaut is moving toward the robot, or the robot needs to travel through the location where the astronaut is lingering). If the conflict is actual, CARIL generates an adaptive response that moves the robot to an out-of-the-way location, delays a planned moved until the 'blocking' astronaut is out of the way, etc.. Importantly, it does this without generating a WH of the cause of the conflict, since it can not act on the WH anyway.

Adding human-robot communication allows the presumed intention in the working hypothesis to be confirmed or disconfirmed with the astronaut involved. However, it also requires both the conflict and the working hypothesis to be represented explicitly in the context representation so they can be used to support the reasoning associated with the ensuing dialog. Similarly, the concept of a dialog itself has to be created and represented, to capture the pragmatic purpose of the dialog as confirming the assumptions in the working hypothesis associated with the underlying conflict.

The individual communications that make up a dialog also had to be explicitly represented and reasoned about, during the span of the dialog. There are at least three reasons for this, all tied to the idea of generality. First, a dialog must be able to be started by an astronaut as well as a robot, so each individual astronaut communication must be viewed as potentially the start of a new astronaut-initiated communication as well as part of an on-going dialog. Second, for robot-initiated dialogs, the concept of a dialog will be created before there are any actual communications. Third, the kind of semantic information that is relevant within a communication will depend on the conflict and working hypothesis that gave cause to the dialog in the first place, so the way in which the individual communication is semantically processed will depend on this context information.

The current CARIL includes generalized sets of knowledge elements to address each of the four classes of potential conflicts listed above. Although space precludes more detailed discussion of all of them, one portion of the reasoning and representational additions is examined in more detail below. The specific case considers the reasoning and actions needed to develop a situationally-appropriate dialog to confirm an astronaut's intentions when moving unexpectedly in the workspace.

Occasionally, the robot may also be scheduled to work collaboratively with astronauts, typically in equipment set-up and put-away around experiments or repair activities. In these collaborations, the action compliance can be more complex, and focuses on understanding when temporal and/or spatial relationships between tasks can be violated, leading to a situation where the robot should adapt its behavior to the situational context. The reasoning to deal with these contexts is an extension of the above approach, and necessarily involves communication with an astronaut to confirm the source of the conflict and/or to negotiate and communicate the robot's adaptive response to it.

A. Detailed View of Confirming a Presumed Intent

This line of reasoning begins with the formation of a WH based on an expected Conflict caused by an astronaut who is remaining in place, off-schedule. The WH explicitly includes a specific presumed intention (e.g, staying late to complete a scheduled task). The Conflict and the WH are represented as explicit concepts on the Conflicts & Adaptation level of the Situational model.

CONFLICT: A conflict concept is formed and posted on the situational representation when the robot is nearing a scheduled movement into another module and realizes that an astronaut has not vacated that module, per that astronaut's schedule. This conflict is represented using the following semantic frame:

Robot<r> IS IN module <loc1> DOING <Activity> NOW

AND

Robot <r> IS scheduled in robot-work-plan [TO <move-to> OR TO <move-through>] MODULE <loc2> AT TIME <NOW-delta>

AND

Astronaut <Y> has been in module < loc2> for <t1> after scheduled departure, as of time <t2>, blocking R's scheduled movement.

Context-based Reasoning: CARIL forms a working-hypothesis (WH) concept of the following form, as a plausible strategy for resolving the conflict:

<u>Working Hypothesis (WH):</u>

Astronaut <y> HAS INTENT TO "stay late at last work activity" AND TO "depart for next scheduled task (<next-Astronaut-Task>)" IN Module <loc3> WITHIN <epsilon> minutes" of NOW

If communication with the astronaut is permitted by the communication strategy (which is also explicitly represented as a policy statement with in the robot's representation of itself on the situational panel), then CARIL initiates a dialog with the astronaut involved to confirm its presumed intention and departure time/destination. The dialog begins by invoking the Confirm-Intent communicative action.

Eventually, the Communication Handler (part of the Action Controller) will be able to choose among multiple Communicative Strategies and to use variable ways to form the message/utterance (to appear less machine-like). Initially, however, the Communication Handler executing the Confirm-Intent action can randomly use one of the following chat messages to initiate an attempt to confirm the Astronaut's intention:

- "Excuse me – I am in <loc1> waiting to do my assigned work in <loc2>. Do you plan to move on to your scheduled equipment repair in <loc3> soon?"
- "Excuse me – Do you plan to move on to your <next-Astronaut-Task> in <loc3> soon?"
- "Excuse me – When will I be able to move into <loc2> and perform my scheduled work there?"

The items in brackets are applied from the conflict and WH frames and passed to the Communication Handler from the CARIL context representation.

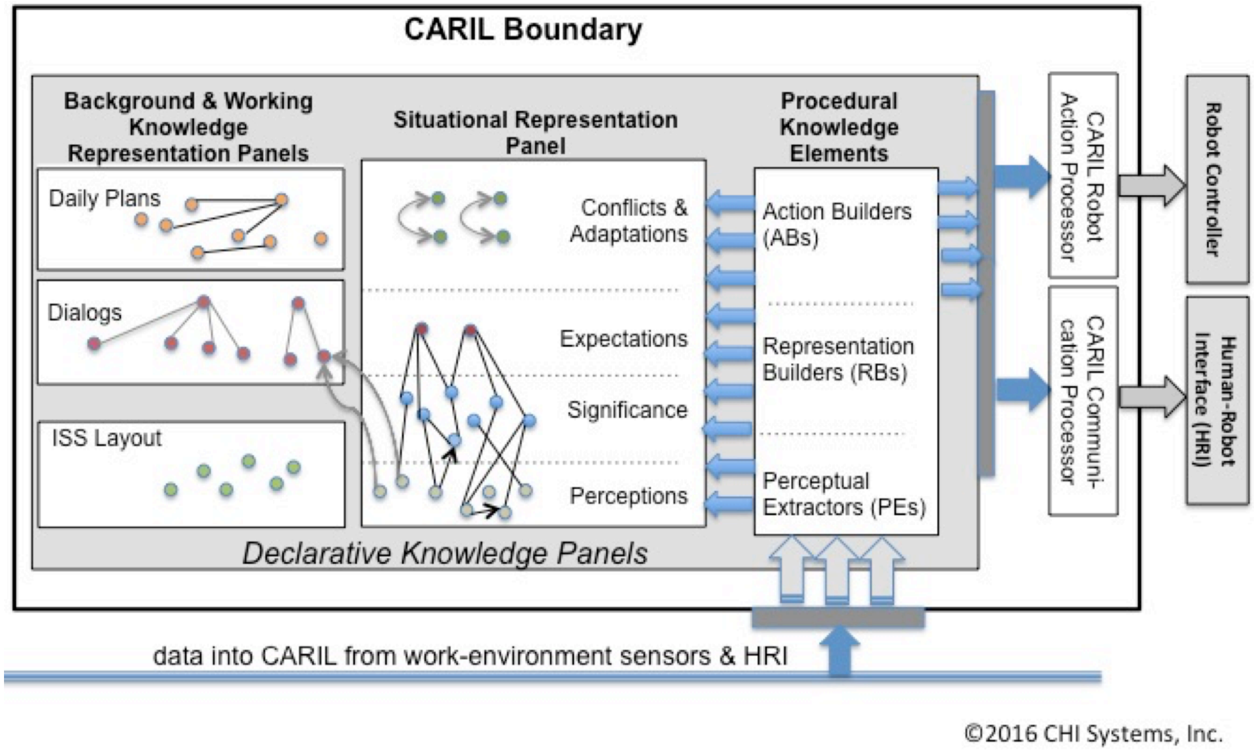

Fig. 1. Organization of knowledge and external interaction processes in CARIL

B. Architectural View of the Space Habitat Application

Figure 1 shows the organization of CARIL. The situational representational panel has a set of concepts at the Adaptations and Conflicts level, specifically to make explicit representations of potential Conflicts and Working Hypotheses as to the root cause and nature of each.

The working knowledge panel for Dialogs is a declarative representation space where CARIL can construct context-sensitive representations of specific dialog with astronauts. Even when completed, a representation of the dialog remains on this panel. This is because past dialogs can provide context for future ones, and thus the persistence of past dialogs allows this aspect of context to be considered in reasoning about current dialogs. Each dialog in this panel has semantic links between it and the various concepts in the Situational representation panel that fill in the context-sensitive slots in that dialog.

The set of Representation Builder knowledge elements contain the knowledge needed to build and maintain the situational representation, including potential conflicts and working hypotheses about conflicts.

The set of Adaptive Action Builder knowledge elements contains the procedural knowledge needed to enable the construction and management of action compliance plans, and of robot human dialogs of various types. Finally, Action controller includes two components, one of which manages physical robot actions by sending action directives to the physical robot controller, and one of which manages robot-human communications, by interacting with a chat-communication mechanism.

C. The ICE Engine

The situational representation and reasoning in CARIL are implemented in general purpose situational and context modeling engine called ICE (the Integrated Context Engine). Thus, CARIL can be considered to be a specific application of the more general ICE software.

Figure 2 show the ICE Architecture. Proper description of the ICE infrastructure requires we that start at the bottom and work back up through the complexity. At the lowest level of ICE, the data storage component, is a combination of Eclipse RDF4J (Version 2.0; 2016) and a customized version of the KiWi Triplestore from the Apache Marmotta project (Version 3.3.0; 2016). RDF4J is a Java-based framework for handling Resource Description Format, or RDF, data while staying agnostic to the specifics of the storage mechanism. KiWi Triplestore, on the other hand, is a SQL-based data backing for RDF4J, providing an efficient method for storing RDF statements in a relational database. In addition to a large performance increase realized through the use of SQL, KiWi Triplestore also provides a more robust ability to monitor transactions resulting in the addition and removal of statements.

Sitting atop the data storage infrastructure is the Pattern Recognition in Motion, or PRIM, Store. The PRIM Store serves two purposes. First, it abstracts away the configuration of the RDF4J / KiWi Triplestore, ensuring all the required components are included in the RDF4J stack. Second, as the name suggests, the PRIM Store is responsible for supporting pattern matching on the RDF data stored in the KiWi Triplestore. Through the PRIM Store, patterns in the format of either partial- statements, i.e. specifying zero or more statement

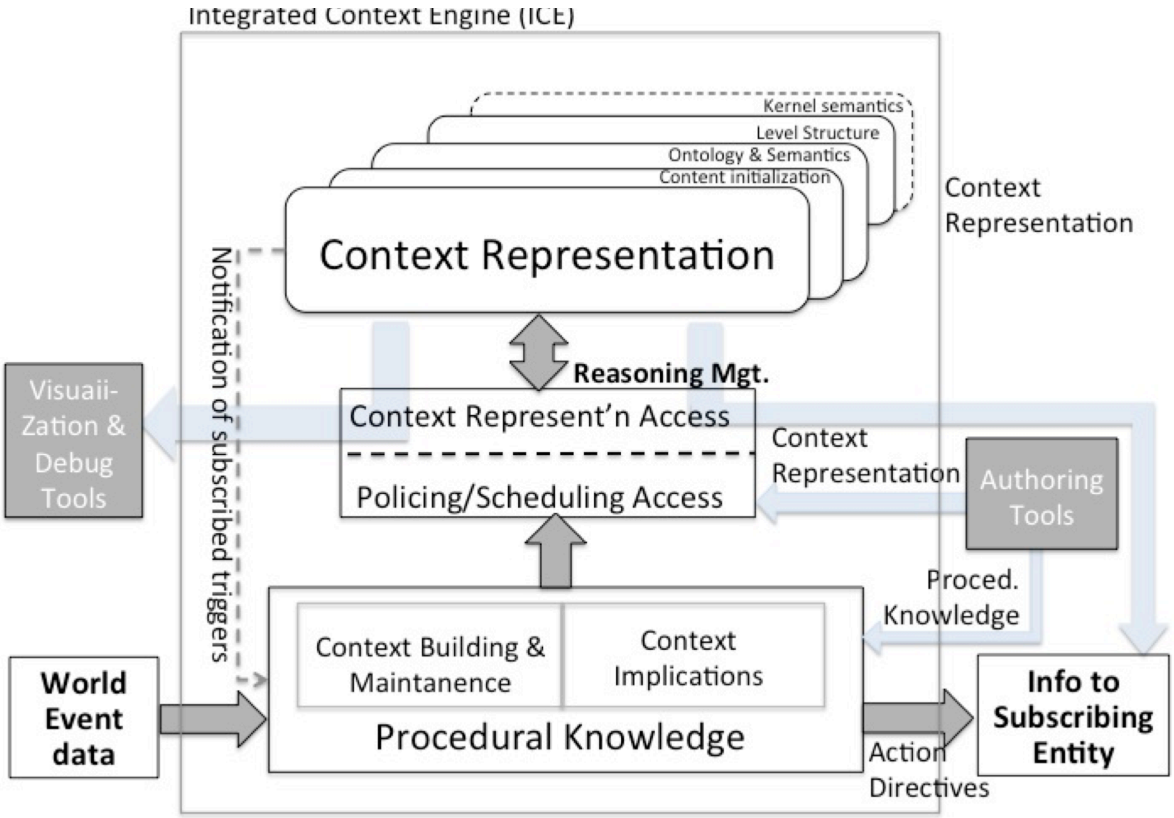

Fig. 2. ICE Architecture showing external interfaces and associated development tools

components, or SPARQL (the RDF equivalent of SQL) queries can be registered and related to callback instructions. Registered patterns are stored in the triplestore itself, allowing for both reasoning over the patterns themselves as well as rebuilding patterns after the datastore has been shut down. Pattern recognition is achieved using many short-lived threads that check each pattern against the changed data before the transaction is made permanent, ensuring that no matches are missed due to rapid changes. In addition to recognizing new pattern matches, it is also possible to monitor when matches are no longer true.

Two PRIM Stores are combined to create the main ICE system. One PRIM Store, the data- PRIM, is used for storing the situation information ingested and fused from the environment, while the second PRIM Store, the meta-PRIM, is used for meta-cognitive information such as information about Knowledge Elements (KEs) registered with the system. The meta-PRIM also stores information detailing the contexts of KEs that are waiting to run, running, and previously finished. The storage of this information allows for detailed visualization of the data flow through KEs. In addition to KEs, ICE also allows the registering of external listeners, components which receive callbacks for pattern matches but are not governed by the ICE control system.

CARIL is a particular configuration of ICE in the situation awareness space involving robots and astronauts in a space habitat modeled on the International Space Station. This configuration encompasses the KEs previously described, as well as the existential information loaded into the system at runtime and the level structure used for organizing information.

The effectiveness of CARIL has been evaluated through the use of an ISS simulator made in JMonkeyEngine (Version 3.0). The simulator provided the ISS environment and simulated astronauts that were used to evaluate the logic of CARIL through scenarios involving non- interference action compliance, including communication. The sensory input provided by the simulator consisted of positional data, time data, and communications. The action affordances consisted of movement, manipulation of objects, and communication.

IV. Conclusion and Next Steps

The main goal of this research was to demonstrate that the CARIL context-based approach and architecture could generate a range of action compliant behaviors. This was done using computer simulation. A full declarative context representation for human-robot work in the space habitat was defined, coded and inserted into the CARIL architecture, as were sets of procedural perceptual, representation-building, and adaptation-building knowledge elements. The CARIL knowledge base included the ability to resolve conflicts either with or without robot-human communication based on either a situational decision, or or based on a general policy (e.g., 'always communicate' or 'never communicate'). In parallel to the knowledge engineering, a simulation was created of the space habitat and of four astronauts performing a range of tasks in the habitat.[1] The astronauts were programmed to deviate in both specific and random ways from the general work plan in specific simulation runs. In each run, CARIL directed the robot in both carrying out its work tasks (which could be independent of, or collaborative with, astronauts' tasks). The simulation was used first to debug, and then to test and demonstrate CARIL's capabilities. As of this writing, CARIL was able to identify and adapt to potential conflicts arising from:

[1] The simulation was created by the Institute for Human-Machine Cognition, by Dr. Matthew Johnson, Dr. Robert Hoffman, and Mr. Daniel Duran.

- the time-based movement cases described above, such as astronaut staying late, etc.;
- unavailability of resources needed for cooperative robot-human tasks, such as when a tool that a robot needs to find and bring to an astronaut is unavailable, and a substitute must to be negotiated, via context-specific communications, with the astronaut; and
- failed temporal coordination constraints, such as when a human task that needed to be completed before a robot could start its next task is abantonded in an incomplete state by an astronaut, requiring the robot to initiate communications to clarify whether or not its task could be started.

Based on these simulations, we have demonstrated that CARIL is capable of generating context-sensitive action compliance in a CARIL-controlled robot.

In future work, we will progress from a simulated environment to a physical environment, using a real robot and real humans, and the same CARIL implementation from the simulation environment to assess and measure the CARIL capabilities and limitations. An important new variable in this work will be employing real (not simulated) robot sensation/perception devices and real robot controllers to translate CARIL directives into action-compliant behaviors.

This will involve developments in three steps. First, we plan on replacing the JMonkey simulation with a Robot Operating System (ROS) [24] environment and simulator to provide a standardized simulation space. Next, we plan on progressing past simulators and evaluating CARIL effectiveness in controlling real robots, namely TurtleBots. And finally, in both the simulated and real-world environments we plan to introduce scenarios involving multiple robots to demonstrate the powerful centralized capability of CARIL.

REFERENCES

[1] R. Pew, and A. Mavor, Eds., *Modeling Human and Organizational Behavior: Application to Military Simulations*, Wash., DC: National Academy Press, 1998.

[2] S. B.Van Hemel, J., MacMillan, and G. L. Zacharias, Eds., *Behavioral Modeling and Simulation: From Individuals to Societies*, Washington, DC: National Academies Press, 2008.

[3] *K. Gluck, and R. Pew, R., Eds., Modeling Human Behavior with Integrated Cognitive Architectures: Comparison, Evaluation, and Validation*, Mahwah, NJ: Erlbaum, 2004.

[4] M. Johnson, J. Bradshaw, P. Feltovich, C.M. Jonker, C. M., M.B. van Riemsdijk, and M. Sierhuis, "Coactive design: designing support for interdependence in joint activity," Journal of Human Robot-Interaction., Vol.3, No. 1, pp. 43-69, 2014.

[5] M.Barnes and F. Jentsch, Eds., *Human-robot-interactions in future military operations.* Farnham: Ashgate, 2010.

[6] T. Kaupp, and A. Makarenko, "Measuring human-robot team effectiveness to determine an appropriate autonomy level," in *Proc. IEEE International Conference on Robotics and Automation*, 2008, pp. 2146-2151.

[7] G. Hoffman, and C. Breazeal, "Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team," in *Proc. of the ACM/IEEE Int. Conf. on Human-Robot Interaction,* 2007, pp. 1-8.

[8] J. Wang, and M. Lewis, "Human control for cooperating robot teams," in *Proc. of the 2nd ACM/IEEE International Conference on Human-Robot Interaction*, 2007, pp. 9-16.

[9] W. W. Zachary, and J.-C. Le Mentec, "Modeling and simulating cooperation and teamwork," Military, Government, and Aerospace Simulation , Vol 32, No. 3 pp.145-150, 2000.

[10] R. Stone, and R. Lavine, Eds., "The social ife of robots," Science, Vol. 346, pp.179-203, 2014.

[11] W. Zachary, W. Weiland, J. Stokes, J. Scolaro, and T. Santarelli, "Using synthetic naturalistic worlds to train teamwork and cooperation," in *Scaled Worlds,* S. Schliffit, L. Elliott, E. Salas, M. Coovert, Eds. Surrey, UK: Ashgate Publishing, 2004, pp. 316-352.

[12] N. J. Cooke, P. A. Kiekel, E. Salas, R. Stout, C. Bowers, and J. Cannon-Bowers, "Measuring team knowledge: A window to the cognitive underpinnings of team performance," Group Dynamics: Theory, Research, and Practice, Vol. 7, No. 3., 2003.

[13] L.A. Suchman, *Plans and situated action: The problem of human-machine communication*, New York, NY: Cambridge University Press, 1987.

[14] R.R. Hoffman, and D. D. Woods, "Beyond Simon's slice: Five fundamental tradeoffs that bound the performance of macrocognitive work systems," IEEE Intelligent Systems, Vol 26, No.6, pp.67-71, 2011.

[15] W. Zachary, A. Rosoff, L. Miller, and S. Read, "Context as cognitive process," in *Proc. of 2013 Semantic Technologies in Intelligence, Defense and Security (STIDS) Confeence*, K. Laskey, I. Emmons, and P. Costa, Eds. Vol.1097, 2013, Avilalble http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol 1097/STIDS2013_T07_ZacharyEtAl.pdf

[16] W. Zachary, M. Johnson, R. Hoffman, T. Thomas, A. Rosoff, and T. Santarelli, "A context-based approach to robot-human interaction," Procedia Manufacturing, Vol. 3, pp. 1052-1059. 2015.

[17] W. Zachary, A. Rosoff, T. Thomas, V. Iordanov, S. Read, and L. Miller, "Toward computational, embeddable, cognitive models of context," Procedia Manufacturing Vol. 3, pp.5269-5276, 2015.

[18] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," Human Factors, Vol. 37, No.1, pp. 32-64, 1995.

[19] M.R. Endsley, "The role of situation awareness in naturalistic decision making," in *Naturalistic decision making*, Mahwah, NJ, Lawrece Erlbaum Associates, 1997, pp. A 97-39376 10-53.

[20] O.G. Selfridge, "Pandemonium: a paradigm for learning," in *Proc. of a Symposium Held at the National Physical Laboratory,* 1958, pp. 513-526.

[21] J. R. Searle, *Speech Act Theory,* Cambridge: Cambridge UP, 1969.

[22] J.R. Searle, F. Kiefer, and M.Bierwisch, Eds. *Speech act theory and pragmatics*, Dordrecht, Holland: D. Reidel Publishing Company, 1980.

[23] J.A. Thomas, *Meaning in interaction: An introduction to pragmatics,* New York, NY: Routledge, 2014.

[24] Robot Operating System [Computer software]. Retrieved from http://www.ros.org. 2016