Envisioning Cognitive Robots for Future Space Exploration

Terry Huntsberger*, Adrian Stoica
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA, USA 91109

ABSTRACT

Cognitive robots in the context of space exploration are envisioned with advanced capabilities of model building, continuous planning/re-planning, self-diagnosis, as well as the ability to exhibit a level of 'understanding' of new situations. An overview of some JPL components (e.g. CASPER, CAMPOUT) and a description of the architecture CARACaS (Control Architecture for Robotic Agent Command and Sensing) that combines these in the context of a cognitive robotic system operating in a various scenarios are presented. Finally, two examples of typical scenarios of a multi-robot construction mission and a human-robot mission, involving direct collaboration with humans is given.

Keywords: Cognitive robotics, behavior-based control, space robotics, hybrid architectures

1. INTRODUCTION

There are a number of different approaches to building a cognitive architecture for control of single and multiple physical robots. Among the cognitive architectures that have been fielded thus far are symbolic/production systems such as Soar[1] and Robo-Soar[2], ACT*[3] and extensions such as ACT-R[4] and ACT-R/E[5], EPIC[6], and ADAPT[7]; connectionist systems such as CTRNN[8], and ART[9] and variants such as ARTMAP[10,11], and Psi[12] and MicroPsi[13]; and hybrid systems[14]. The symbolic processing systems are derived from studies of human cognition, and as such don't always port well onto robotic platforms, particularly in the area of multi-robot cooperation. ACT-R/E addresses this issue through the addition of an internal simulation process to ACT-R that mimics the 'like-me' behavior evidenced in primates and humans[15]. This simulation process enables the cognitive system to map external sensed actions by other agents into its own behavior base (walking in someone else's shoes, so to speak). There have been numerous studies into the common ground between cognitive processing and formal process algebras[38-40]. JPL has developed a tightly integrated instantiation of a cognitive agent called CARACaS (Control Architecture for Robotic Agent Command and Sensing), a block diagram of which is shown in Figure 1, to address many of the issues for survivable, autonomous unmanned vehicle control[23-25]. CARACaS is composed of a Dynamic Planning Engine (currently CASPER), a Behavior Engine (currently

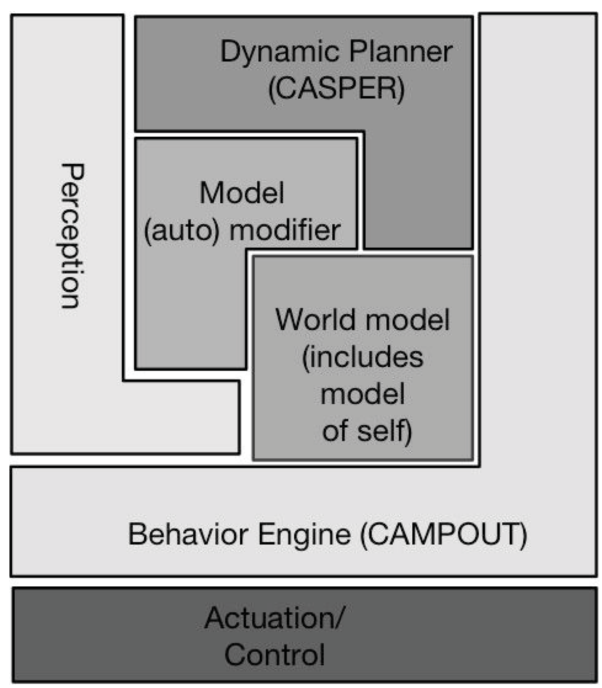

Figure 1. Block diagram of the CARACaS system. Connections between components are indicated by relative proximity.

* terry.huntsberger@jpl.nasa.gov, ph: 818-354-5794; fax: 818-393-5007.

CAMPOUT), a Perception Engine, and a World Model. The Behavior Engine in CARACaS is mapped into a process algebra formalism in order to maintain a linear complexity for inferring sensed behaviors of other agents.

This internal simulation process also addresses another key problem that needs to be solved by any "intelligent" system regardless of its level of consciousness, that of being able to operate in new, previously "unseen" contexts - that are not simple small deviations from previously seen contexts, but dramatically different[16]. Such anticipatory capabilities have been studied in the context of organic and robotic systems[17-20]. While small deviations may be handled in a simple interpolation/extrapolation framework (neural, fuzzy, etc), a totally novel context may 'confuse' a robot. Such contexts could be sensing artifacts (e.g. deteriorated sensor/vision chip in which all images get alternating black bars of certain thickness, rendering embedded object recognition routines useless, to seeing objects of very different shape, color and variation in time compared to previously seen/memorized ones). These unseen contexts may require actions that are novel combinations of action primitives. Yet in order to determine how appropriate these combinations are, it is useful to exercise a capability to determine the answer to "what if" type scenarios.

Running an internal simulation with potential robot actions as inputs to a model of the world will generate consequences that are to be evaluated, ranked, and from which optimal actions can be selected. Other "what if" input context may not be initiated by the robot, but by other actors in the world model (e.g. what if I lost one sensory modality, what if another robot comes towards me at high speed, etc). The results of such simulations could be used not only for determining robot behavior, but also for improving its model of the world and raising its own level of capability in the presence of a teacher. In the worst case, the teacher generates only reinforcement (self-guided experimentation may generate similar effects, albeit at higher risks and energy consumption), while in a better case the teacher could suggest "what if" contexts and indicate at least preferences if not exact responses for such contexts. Particularly interesting in the context of implementations of simulation theories is the work by Svensson[21], which offers the Representation-as-Simulation Hypothesis (RaSH) thesis, seeing simulation processes as off-line representations.

A recent study that evaluated issues associated with remote interaction with an autonomous vehicle within the framework of grounding found that missing contextual information was a recurring problem for the operations team[22]. This missing contextual information led to uncertainty in interpretation of data that was collected and possible errors in how the autonomous vehicle was commanded. The problems that were encountered increased as the remote agent became more and more autonomous through activation of additional capabilities. Behavior of the remotely located autonomous vehicle would not always fit the "mental model" of the operators, leading to inefficient use of the platform.

One of the conclusions of the study was that the common ground would be better established if the autonomous agents could describe what they do and why. This capability is provided if the robotic agents have enough onboard self-awareness to dynamically adjust the information conveyed back to the operator based on a detail level component analysis of requests. A cognitive system that provides a <u>formal mathematical basis</u> for <u>onboard representation of the behavior-based control</u> of autonomous agents, combined with an <u>integrated, adaptive explanation capability</u> can provide <u>common grounding</u> between operator and vehicle.

Fielding robots in space places some restraints on what can be done due to limited computing capabilities, mass and volume constraints, low bandwidth communication channels coupled with oftentimes long delays, and power issues. Several key aspects of a cognitive approach to unmanned vehicle control for space exploration include the handling of the inherently uncertain nature of dynamic surface operations, sensing for hazard detection/avoidance and situation awareness, behaviors for obeying the social rules during interactions with other manned and unmanned vehicles, cooperation among heterogeneous vehicles, onboard resource-based planning for mission operations, integrated system health maintenance for long duration missions, and the human operator command interface. The processes running within CARACaS include *reactive* processes for autonomous safe navigation and path planning, *deliberative* processes for planning and reasoning about complex, possibly conflicting goals during mission operations, and *reflective* processes for resource management and self-preservation.

Reactive components in space robotics require deterministic reaction to unanticipated occurrences which can be captured through three components: (1) reacting to the occurrence with an appropriate response, (2) reacting to the occurrence within a predictable timeframe, and (3) providing other system components with updated autonomous vehicle state information. The first requirement is met by using a behavior coordination mechanism based on Multiple Objective Decision Theory (MODT) that guarantees a solution that is "good enough" within mission constraints. The second requirement is met with finite state machines using embedded resource and timing operators to define the tactical behavior network. The third requirement is met by the direct feedback loop from the behavior network to the reasoning/planning components of the system for internal state information transfer using a common shared format.

The next section presents the individual components of CARACaS, followed by a brief discussion of the formal methods used for the implementation in terms of process algebras. This is followed by application of the system to two scenarios, and finally closing with a summary, discussion of results and references.

2. COGNITIVE SYSTEM – OVERALL ORGANIZATION

Dynamic Planner

The Dynamic Planner leverages the CASPER (Continuous Activity Scheduling Planning Execution and Replanning) continuous planner[26] developed at JPL. Given an input set of mission goals and the autonomous vehicle's current state, CASPER generates a plan of activities that satisfies as many goals as possible while still obeying relevant resource constraints and operation rules. The "what-if" capabilities of CARACaS are based on the event horizon look-ahead view that CASPER maintains throughout the mission. CASPER has been used to autonomously perform the planning/re-planning for the Earth Observation 1 (EO1) satellite[27] continuously since November 2004 and recently won the NASA Software of the Year Award. A description of the autonomous vehicle, including resources and state information, as well as applicable mission and operations rules is encoded in the planner's modeling language. Plans are dynamically updated using an iterative repair algorithm that classifies plan conflicts (such as a resource over-subscription) and resolves them individually by performing one or more plan modifications. CARACaS takes a most-committed, local, heuristic, *iterative repair approach* to producing and modifying plans. This approach gives CARACaS the advantages of 1) allowing the repair algorithm to be applied at any time and on any given plan (abstract or detailed), 2) enabling fast re-planning when conditions or goals change, 3) allowing the easy incorporation of heuristics to prune the search space, and 4) incurring less overhead during search since a local repair algorithm does not require the saving of intermediate plans or backtracking points.

Behavior Engine

CARACaS leverages the results of previous efforts at JPL in the multi-agent control architecture CAMPOUT (Control Architecture for Multi-robot Planetary Outposts)[28-32] in order to develop behavior composition and coordination mechanisms. CARACaS uses *finite state machines* for composition of the behavior network for any given mission scenarios. These finite state machines give it the capability of producing formally correct behavior kernels that guarantee predictable performance using formal methods (Labeled Transition Systems, see next Section).

For the behavior coordination mechanism (BCM) CARACaS uses a method based on *Multi-Objective Decision Theory (MODT)* that combines recommendations from multiple behaviors to form a set of control actions that represents their consensus. This approach provides for a coordination scheme that allows all behaviors to simultaneously contribute to the control of the system in a cooperative rather than a competitive manner, which explicitly addresses tasks that may have conflicting goals. CARACaS uses the MODT framework[33] coupled with the interval criterion weights method[34,35] to systematically narrow down the set of possible solutions (size of space grows exponentially with the number of actions), producing an output within a time-span that is orders of magnitude faster than a brute force search of the action space.

Perception

The Perception Engine leverages algorithms derived from those used onboard the Mars Exploration Rovers (MER) for passive stereo imaging, hazard detection, and visual localization for navigation. Camera models based on polynomial expansions used to correct camera/lens distortions are derived from a series of images obtained during a calibration procedure. A fast stereo algorithm developed at JPL[36] is used to generate a range map for hazard avoidance and sensing of other agents during the motion.

World Model

The World Model in CARACaS is based on explicit state knowledge of the robot and other agents that are in the same environment. The state knowledge of other agents is a mixture of information communicated from the other agents directly depending on bandwidth or through onboard sensing, as well as the anticipated states that are known from the mission plan. In addition, there are short-term-memory and long-term-memory components that interface to the behavior engine[37], the dynamic planner and the perception submodules. There is also a global map of the mission area derived from any map information previously obtained (satellite, etc) supplemented with local sensory inputs from the agent and other agents that are in the area.

The symbolic processing aspects of systems like SOAR are captured in CARACaS through the mapping of the behaviors into a process algebra that provides formal symbolic statement composition and inference operators. The inference operation for the 'like-me' analysis occurs when a sensory measurement is made and compared to the existing behavior base. SOAR is limited to first-order logic operations, whereas the Cost Calculus[41] includes an explicit representation of unknown or uncertain information in the logic operations. The connectionist components of CARACaS are contained in the short term memory of the system, so in some sense CARACaS is a hybrid system.

3. COGNITIVE SYSTEM - FORMAL BASIS

Formal process algebras were originally developed to model and analyze distributed computation and communication processes. These types of processes have common features with the cognitive robotics community such as temporal sequencing, uncertainty representation, self/multi-system awareness, symbolic processing, and perceived cost. Features such as feelings and emotions don't map directly into the process algebras, but can be phrased respectively as internal state awareness and action-urgency.

3.1 Cost Calculus

A Cost Calculus ($-Calculus)[41] is a model for resource bounded computation based on process algebras that:

1. Provides a means for generating incremental solutions for computationally hard, real-life problems
2. Provides a uniform representation for the use of uncertain information during the cost-optimization process (kΩ-optimization)
3. Provides an explicit representation of unobservable behavior (incomplete knowledge about an agent or the environment) using the silent (invisible) action ε - particularly important since most sensors have limited range and will not be able to provide all needed information for decision making all of the time
4. Currently is the basis for the CCL (Common Control Language) developed under ONR funding used for control of UUV (Unmanned Undersea Vehicles) at the Naval Undersea Warfare Center, Newport[42,43].

Behaviors are written as $-expressions organized into 6 sets (maneuver, navigate, communicate, configure, monitor/report, and execute convention). $-expressions are built using the algebraic operators of send/receive, cost assignment, defined simple/process call and sequential/parallel composition.

The Behavior Engine in CARACaS is mapped to a Cost-Calculus ($-Calculus) framework[41], and observed behaviors from other agents are matched to the existing tactical behavior base using well-known bisimulation equivalence relations. Bisimulation equivalences are binary relations between state transition systems, associating systems that behave in the same way in the sense that one system simulates the other and vice-versa. The advantages of this approach with regard to the current state of practice:

1. Easy to integrate for testing on autonomous vehicles due to the existing base of tested behaviors already running onboard technology rovers under CARACaS
2. Reduced computational complexity with respect to existing algorithms (linear vs. polynomial or exponential) leading to efficient onboard use
3. Rigorous mathematical foundation (Process Algebras) that supports analysis
4. First approach to explicitly factor in sensing from a moving platform and analysis of actions by other independent agents in the surrounding environment

This approach has been used successfully in a number of different fields, including motor schemas for robotic control and plan recognition in economic processes.

3.2 Inference

For the inference of sensed behaviors, the observation equivalence of behaviors on a single autonomous agent and between two or more agents is done through bisimulation relations:

1. Behaviors are formally expressed as a LTS (Labeled Transition System)[44], and as such, the bisimulation equivalence can be established in linear time[45]
2. Linear efficiency enables both onboard and/or offboard use of the technique

3.3 Learning

For the learning of sensed behaviors that is necessary for common grounding of behavior sequences that were not previously observed or in the command dictionary of the autonomous agent, reinforcement learning of observed behavior patterns is used. Reinforcement learning is a mode of the $k\Omega$ optimization built into the $-Calculus and:

1. LTS representation used for generation of history of behavior use (similar to networks of Michaud & Mataric[46])
2. Q-learning update equation is directly represented in the $-Calculus using cost / general choice and sequential composition operators

Efficient onboard reinforcement learning algorithms have been previously developed at JPL and demonstrated for adaptive behavior in out-door environments on rovers[47,48]. These experiments used "rover health" defined in terms of available power and goal achievement as the objective function for the learning. The Q-learning component was not used because the behavior base is currently relatively small and the sequences built using the sequential composition operators could be exhaustible parsed.

3.4 Explanation capabilities

For the development of explanation capabilities, a dynamic decision tree decomposition[49] of the observed behaviors is used to generate a set of rules for explanation:

1. Decision tree generation uses information gain and pruning to limit the size of the tree
2. Rules are evaluated based on hit rates, miss rates, pessimistic error rate, and information gain

An adaptive level of detail is automatically built into this process in that all of the sensory information that led to a behavior is available and can be conveyed to the operator if the Human Machine Interface (HMI) has a detail level of request capability.

3.5 Cognitive skill rating

ConsScale (*Cons*ciousness *Scale*) levels introduced in Arrabales, et al.[16] are used for a qualitative assessment of the cognitive skills of our system characterized by the architecture in Figure 1. The ConsScale levels range from 1 to 11, with 1 (Decontrolled) having no relationships defined between sensors and actions, and 11 (Super-Conscious) having the ability to synchronize and coordinate multiple streams of consciousness. In accordance with the authors' indication that the metric needs to be seen in the context of the application specific domain, our architecture would be at a ConsScale level 6 (Emotional). This includes all levels below, and matches specific Cognitive Skills (CS) detailed in Table I in [Arrabales, et al., 2009] to Level 6: $CS_{6.1}$ Self-status assessment (which we do not however interpret as background emotions); $CS_{6.2}$: Status assessment (background emotions) cause effects in agent's body; $CS_{6.3}$ Representation of the effect of emotions in the organism; $CS_{6.4}$: Ability to maintain a precise and updated map of body schema; $CS_{6.5}$: Abstract learning. The primary difference between human and robotic cognition lies in this redefinition of "emotion" as internal state for the robotic system. The CLS (cumulative level score) of CARACaS is approximately 1.68, while the CQS (cognitive quantitative score) of our system is 101.08 (for comparison, self-consciousness is 200, super-conscious at 1000). The calculation for CSL and CSQ come from equivalent of full level 6.

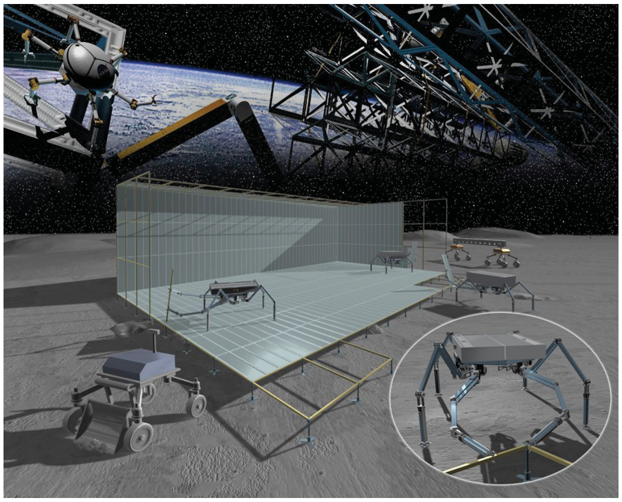

Figure 2: Artist's concept of a variety of surface and on-orbit assembly and construction operations, including truss assembly, component transport, and site preparation. A heterogeneous team of robotic agents are shown, with both wheeled and walking

Figure 3. Coordinated transport of extended container (2.5 meters) by SRR and SRR2K, as performed in Arroyo Seco near JPL. (Left) row transport formation; (Right): column (leader-follower) transport formation

4. SCENARIOS

4.1 Scenario 1: Multi-robot construction

Objects that are four to five times the length of a single mobile platform are extremely difficult to manipulate and transport. The *Robot Work Crew (RWC)* concept assumes use of multiple rovers for coordinated operations such as those shown in Figure 2. Operations performed on such an extended payload include using two rovers that are cooperating to carry the beam over uneven terrain, with examples of *row* and *column* transport being shown in Figure 3. CAMPOUT, the behavior-based core of CARACaS, included goal arbitration in the outdoor environment in determining whether to execute *Go-to-Goal*, *Avoid Obstacles*, and *Reconfigure Payload*. These goals were internally represented using a Parallel Composition operator under the $-Calculus. The goal for the experimental study was the transport of an extended container by two rovers (SRR and SRR2K, the latter being a minimalist mechanization of the first) from a pickup point to a deployment zone that is up to 50 meters away, over unoccluded natural terrain. This was accomplished with a four-phase sequence that involved numerous realignments between the rovers due to load shifting when going over uneven ground. The rovers worked as a team with self-awareness of their roles in the transport process. The rovers anticipated when to make the corrections based on stress loads along the shared beam.

As a general strategy, explicit communication between the rovers was minimized, as reflects possible operational constraints during an actual mission. This tailoring of the communication is facilitated by using the shared container as an implicit means of communication—e.g., relative positions of the rovers are known through the yaw gimbal angle on each rover. Also, we are exploiting natural design constraints of the task where possible to assess useful trades of mechanized cooperation versus explicit closed loop controls (as one example, the use of passive compliance in both grippers along the beam axis). The number of simultaneous goals was kept low (3) in order to emulate a system constrained by limited computational capabilities. The current World Model in CARACaS is not a rich enough representation to be considered fully conscious for Scenario 1, being limited to short-term knowledge of the higher level goals of the construction process beyond the assembly manual for the structure. Full consciousness would include more context knowledge of the purposes of individual sub-components and sub-structures within the construction site.

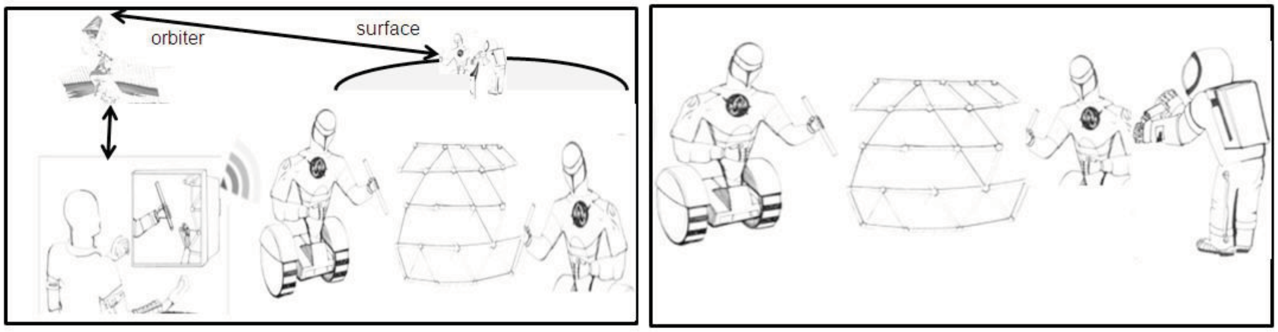

Figure 4. Sketch illustrating two cases in Scenario 3. in which astronauts, either in orbiting spacecraft (left drawing), or on the planetary/lunar surface (right drawing), collaborate with robots, for various operations, such as repair, assembly, etc. .

4.2 Scenario 2: Manned mission, humans collaborating with robots on the surface

The Scenario focuses on human-robot collaboration, in the context of manned missions and robots operating in-situ on planetary/lunar surface. Astronauts may be interacting from an orbiting spacecraft or directly from surface as illustrated in Figure 4 (or a mixed case in which some astronaut is in orbit and others on the surface).

Being in orbit means that there will be times of direct overhead viewing, and other cases when there would be perhaps an indirect viewing context via other orbiting spacecraft. A better global view is feasible, although at a lower resolution. Local imagery may be transmitted by the robots or other surface infrastructure imposing inherent bandwidth limitations/tradeoffs. A certain degree of teleoperation is possible, the limitation not being communication delays (as it would be if trying to control from Earth) but possibly bandwidth and power restrictions at the robot end. If operation is to continue while satellite is not in view, or during the dark hours, robots need to have a high degree of autonomy.

An important cognitive ability is determining the intent of other entities operating in the environment, which has predictive value. Joint attention (focusing attention to the same target/object that another entity/human is attentive to) assists in determining intent ('mind reading'), and complements other cues, including various communication means such as language communication. There is a body of literature focuses on joined attention (see for example the discussion in Sumioka[50], and references cited therein). The traditional context is to exploit face expression. However, in the context of astronauts working in EVA suits (extra-vehicular activities) direct face observation by a robot is made difficult by the helmet. In such cases where robots cooperate with astronauts, embodiment of sensors inside the suit that record things such as eye movements (gaze), as well as possibly using biological signals (EEG, EMG, etc), with transmission of the signals to the robot is an alternate approach.

One can conceive that the information about face expression or gaze direction is in fact processed by the EVA suit and broadcast to all participants in the scene, humans or robots, which in this case have to process less information, and also can receive it even if they are in a position in which they could not see the astronaut face directly (e.g. both astronaut and robot looking forward – without the robot needing to continuously alternate/shift gaze from human to object of human attention)

Assume an assembly task. Perception (stereo vision, facilitated by artificial lighting at night) facilitates a situation assessment. The current most important goal is to continue to add beams to a structure under assembly. The Dynamic Planner determines a sequence of behaviors: find the beam in the workspace, retrieve beam, approach assembly structure from direction of next element insertion, insert beam in place, etc. Perception provides recognition of assembly elements (say beams) and their location in the scene, and during the entire process provides the Behavior Engine with context updates, while the Behavior Engine guides perception for retrieving need information, etc. In the context of the higher-level architecture of human-robot collaboration, one can share roles in determining the next area of assembly structure to be completed, negotiate break times for battery recharge or inspection, or even to determine at a low level which beam to place next. In general an optimization at this level is not only feasible but also advantageous, as illustrated for example in Smith [51].

The optimization of task allocation between humans and robots becomes even more critical in the context of astronauts cooperating with robots on the surface. The humans have a direct/unobstructed view of the task area, and could more easily teleoperate the robots. Conventional teleoperation is however inefficient and tiring for astronauts. Novel human-robot interfaces, using biological signals collected by sensors embedded in the suits (such as bio-sleeves collection EMG, bio-caps collecting EEG, etc) would provide friendlier interfaces and higher efficiency in operation, with greater bandwidths. For better efficiency however one should have the astronaut in supervisory or advisory roles to teams of multiple robots. Embedding sensors in the astronaut suit – particularly in the helmet, would allow collection of information about his face expression, including direction of gaze, which would be processed (and combined with verbal and other cues) to infer intent, etc.

The cognitive mechanisms for the assembly task have many aspects in common with multi-robot assembly operations in the previous scenario. Yet, in this case the Planner would take in consideration other aspects including maintaining a safe work environment for the human (a context in which the Behavior Engine would engage safe behaviors of avoiding human proximity, reduction of power level and speed in human proximity) while Perception should, with a high priority obtain information about human position, trajectory, etc. direction of gaze, object manipulate, etc, all these being interpreted and updated in the world model. CARACaS is currently limited in its ability to include human agents, mostly due to the lack of capabilities for inferring human intent directly from sensory input. The behavior base would have to be supplemented in order to emulate “like-me” behavior.

5. SUMMARY

The cognitive architecture CARACaS was presented and some example space-based scenarios were discussed. CARACaS has been tested extensively on Unmanned Surface Vehicles (USV’s) under US Navy contracts as well[23-25]. Cognitive characteristics of self-awareness (each rover knew its role in the team), anticipatory planning (look-ahead projection of the convex hull of the two rover configuration was used to ensure clearance between the ensemble and hazards), and “like-me” behavior (projection along the shared beam was used to project relative orientation and current activities of the other rover) were demonstrated in the field in Scenario 1. Although Scenario 1 can be done using traditional robotic methods, as the complexity of the site increases, higher level consciousness characteristics of the system become important for inferring the state of the rover in order to stay safe in the possibly highly cluttered environment.

CARACaS rated a Level 6 on the qualitative Consciousness Scale of Arrabales[16], meaning that the system has the ability to generalize its learned behaviors and possesses “feelings” (agent well-being in this case). Current directions include the addition of model modification capabilities so that the World Model will be better able to represent alternate views of the environment around the agents. Explanation capabilities are limited to report of states within the sequences, and need a more intuitive grounding within the overall goal of the construction task. Also, the development and testing of alternate command methods such as gestures and teaching by example are currently being investigated.

ACKNOWLEDGMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors would like to thank Dr. Neville Marzwell of the Advanced Programs Office at JPL for the initial support of the work and further funding under Contract N0001410IP20075 from the Office of Naval Research. We would also like to thank members of the Mobility and Robotic Systems Section at JPL for their support in the development of CARACaS.

REFERENCES

[1] Laird, J.E., Newell, A. and Rosenbloom, P.S., "SOAR: An Architecture for General Intelligence," Artificial Intelligence 33, 1-64 (1987).

[2] Laird, J.E., Yager, E.S., Hucka, M. and Tuck, C.M., "Robo-SOAR: An Integration of External Interaction, Planning, and Learning Using SOAR," [Toward Learning Robots], MIT Press, Cambridge, MA, 113-130 (1993).

[3] Anderson, J.R., [The Architecture of Cognition], Harvard University Press, Cambridge, MA, (1983).
[4] Anderson, J.R., Matessa, M. and Lebiere, C., "ACT-R: A Theory of Higher Level Cognition and its Relation to Visual Attention," Human-Computer Interaction 12(4), 439-462 (1997).
[5] Kennedy, W. G., Bugajska, M. D., Harrison, A. M., Trafton, J. G., "'Like-Me' Simulation as an Effective and Cognitively Plausible Basis for Social Robotics," International Journal of Social Robotics 1(2), 181-194 (2009).
[6] Meyer, D.E. and Kieras, D.E., "A Computational Theory of Executive Cognitive Processes and Multiple Task Performance: Part 1. Basic Mechanisms," Psychological Review 104, 3-65 (1997).
[7] Benjamin, D.P., Lyons, D. and Lonsdale, D., "ADAPT: A Cognitive Architecture for Robotics," Proc. 2004 International Conference on Cognitive Modeling (ICCM), (2004).
[8] Beer, R. D., "Toward the evolution of dynamical neural networks for minimally cognitive behavior," Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior : From Animals to Animats 4, 421-429 (1996).
[9] Grossberg, S., "The Link Between Brain, Learning, Attention, and Consciousness," Consciousness and Cognition 8, 1-44 (1997).
[10] Carpenter, G., "Distributed Learning, Recognition, and Prediction by ART and ARTMAP Neural Networks," Neural Networks 10, 1473-1494 (1997).
[11] Carpenter, G., "Large-Scale Neural Systems for Vision and Cognition," Proc. International Joint Conference on Neural Networks, 454-459 (2009).
[12] Bach, J., [Principles of Synthetic Intelligence Psi: An Architecture of Motivated Cognition], Oxford University Press, New York & Cary, NC, 241-248 (2009).
[13] Bach, J., Bauer, C. and Vuine, R., "MicroPsi: Contributions to a Broad Architecture of Cognition," Lecture Notes in Computer Science 4314, 7-18 (2007).
[14] [Computational Architectures Integrating Neural and Symbolic Processes: A Perspective on the State of the Art], (ed. R. Sun and L.A. Bookman), Kluwer Academic Publishers, Needham, MA (1994).
[15] Meltzoff, A. N., "'Like Me': A Foundation for Social Cognition," Developmental Science 10(1), 126–134 (2007).
[16] Arrabales, R., Ledezma, A. and Sanchis, A., "Establishing a Roadmap and Metrics for Conscious Machines Development," Proc. 8th Int. Conf. on Cognitive Informatics (ICCI 2009), 94-101 (2009).
[17] Cassimatis, N. L., Trafton, J., G., Bugajska, M. D., & Schultz, A. C., (2004). "Integrating Cognition, Perception and Action through Mental Simulation in Robots," Journal of Robotics and Autonomous Systems 49(1-2), 13-23 (2004).
[18] Pezzulo, G. "Anticipation and Future-Oriented Capabilities in Natural and Artificial Cognition," Lecture Notes in Computer Science 4850, 257-270 (2007).
[19] Morse, A.F., Lowe, R. and Ziemke, T., "A Neurocomputational Model of Anticipation and Sustained Inattentional Blindness in Hierarchies," Lecture Notes in Computer Science 5499, 152-169 (2009).
[20] Stepp, N. and Turvey, M.T., "On Strong Anticipation," Cognitive Systems Research 11(2), 148-164 (2010).
[21] Svensson, H., [Embodied Simulation as Off-Line Representation], Licentiate of Philosophy Thesis, Department of Computer and Information Science, Linköpings Universitet, Linköping, Sweden, (2007).
[22] Stubbs, K., Hinds, P. and Wettergreen, D., "Challenges to Grounding in Human-Robot Collaboration: Errors and Miscommunications in Remote Exploration Robotics," CMU Technical Report, CMU-RI-TR-06-32, (2006).
[23] Huntsberger, T., Aghazarian, H., Castano, A., Woodward, G., Padgett, C., Gaines, D. and Buzzell, C., "Intelligent Autonomy for Unmanned Sea Surface and Underwater Vehicles," Proc. AUVSI North America 2008 Unmanned Vehicles Conference, (2008).
[24] Huntsberger, T., Aghazarian, H., Gaines, D., Garrett, M. and Sibley, G., "Autonomous Operation of Unmanned Surface Vehicles (USV's)," Proc. IEEE ICRA Workshop on Robots in Challenging and Hazardous Environments, (2007).
[25] Hansen, E., Huntsberger, T. and Elkins, L., "Autonomous Maritime Navigation - Developing Autonomy Skill Sets for USVs," Proc. SPIE Defense and Security Symposium 2006, (2006).
[26] Chien, S., Knight, R., Stechert, A., Sherwood, R. and Rabideau, G., "Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling," Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS-2000), (2000).
[27] Chien, S., Sherwood, R., Tran, D., Castano, R., Cichy, B., Davies, A., Rabideau, G., Tang, N., Burl, M., Mandl, D., Frye, S., Hengemihle, J., Agostino, J., Bote, R., Trout, B., Shulman, S., Ungar, S., Van Gaasbeck, J., Boyer, D., Griffin, M., Burke, H., Greeley, R., Doggett, T., Williams, K., Baker, V. and Dohm, J., "Autonomous Science on the EO-1 Mission," Proc. of the Intern. Sympos. on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2003), (2003).

[28] Stroupe, A., Okon, A., Robinson, M., Huntsberger, T., Aghazarian, H. and Baumgartner, E., "Sustainable Cooperative Robotic Technologies for Human and Robotic Outpost Infrastructure Construction and Maintenance," Autonomous Robots 20(2), 113-123 (2006).
[29] Huntsberger, T.L., Trebi-Ollennu, A., Aghazarian, H., Schenker, P.S., Pirjanian, P. and Nayar, H.D., "Distributed Control of Multi-Robot Systems Engaged in Tightly Coupled Tasks," Autonomous Robots 17, 79-92 (2004).
[30] Huntsberger, T., Pirjanian, P., Trebi-Ollennu, A., Nayar, H.D., Aghazarian, H., Ganino, A., Garrett, M., Joshi, S.S and Schenker, P.S., "CAMPOUT: A Control Architecture for Tightly Coupled Coordination of Multi-Robot Systems for Planetary Surface Exploration," IEEE Trans. Systems, Man & Cybernetics, Part A: Systems and Humans, Special Issue on Collective Intelligence 33(5), 550-559 (2003).
[31] Huntsberger, T., Cheng, Y., Baumgartner, E.T., Robinson, M. and Schenker, P.S., "Sensory Fusion for Planetary Surface Robotic Navigation, Rendezvous, and Manipulation Operations," Proc. Intern. Conf. on Advanced Robotics (ICAR'03*)*, 1417-1424 (2003.
[32] Huntsberger, T., Aghazarian, H., Cheng, Y., Baumgartner, E.T., Tunstel, E., Leger, C., Trebi-Ollennu, A. and P.S. Schenker, P.S., "Rover Autonomy for Long Range Navigation and Science Data Acquisition on Planetary Surfaces," Proc. IEEE International Conf. on Robotics and Automation (ICRA2002), 3161-3168 (2002).
[33] Pirjanian, P., "Multiple objective behavior-based control," Journal of Robotics and Autonomous Systems, 31(1-2), 53-60 (2000).
[34] Benjamin, M.R., "Multi-Objective Autonomous Vehicle Navigation in the Presence of Cooperative and Adversarial Moving Contacts," Proceedings of OCEANS 2002, (2002).
[35] Benjamin, M.R., [Interval Programming: A Multi-Objective Optimization Model for Autonomous Vehicle Control], Ph.D. Thesis, Department of Computer Science, Brown University, Providence, RI, (2002).
[36] Matthies, L., "Stereo vision for planetary rovers: Stochastic modeling to near real-time implementation," Int. J. Computer Vision 8, 71-91 (1992).
[37] Huntsberger, T.L. and Rose, J., "BISMARC: A Biologically Inspired System for Map-based Autonomous Rover Control," Neural Networks 11(7/8), 1497-1510 (1998).
[38] Forth, J. and Shanahan, M., "Indirect and Conditional Sensing in the Event Calculus," Proc. 16th European Conference on Artificial Intelligence (ECAI 2004), 900-904 (2004).
[39] Shanahan, M. and Witkowski, M., "Event Calculus Planning Through Satisfiability," Journal of Logic and Computation 14(5), 731-745 (2004).
[40] Shanahan, M.P, "The Event Calculus Explained," Artificial Intelligence Today, (ed. M.J.Wooldridge and M.Veloso), Springer Lecture Notes in Artificial Intelligence 1600, 409-430 (1999).
[41] Eberbach, E., "$-Calculus of Bounded Rational Agents: Flexible Optimization as Search under Bounded Resources in Interactive Systems," Fundamenta Informaticae 68, 47-102 (2005).
[42] C. M. Buzzell, [A Common Control Language for Multiple Autonomous Undersea Vehicle Cooperation], Master's Thesis, University of Massachusetts Dartmouth, (2004).
[43] Duarte, C.N., Martel, G.R., Buzzell, C., Komerska, R., Mapparapu, S., Chapel, S., Blidberg, D.R. and Nitzel, R., "A Common Control Language To Support Multiple Cooperating AUVs," Proceedings of the 14th International Symposium on Unmanned Untethered Submersible Technology, (2005).
[44] Milner, R., "Operational and Algebraic Semantics of Concurrent Processes," in [Handbook of Theoretical Computer Science, Vol. B: Formal Models and Semantics], (ed. J. van Leeuwen), Elsevier and The MIT Press, 1203-1242 (1994).
[45] Paige, R. and Tarjan, R.E., "Three partition refinement algorithms," SIAM Journal on Computing 16(6), 973–989 (1987).
[46] Michaud, F. and Mataric, M.J., "Representation of behavioral history for learning in nonstationary conditions," Robotics and Autonomous Systems 29(2-3), 187-200, (1999).
[47] Huntsberger, T., Aghazarian, H. and Tunstel, E., "Onboard Adaptive Learning for Planetary Surface Rover Control in Rough Terrain," Proc. IEEE International Conference on Robotics and Automation (ICRA2005)*,* (2005).
[48] Huntsberger, T., "Onboard Learning of Adaptive Behavior: Biologically Inspired and Formal Methods," Proc. 2nd Symposium on Learning and Adaptive Behavior in Robotic Systems (LAB-RS2009), (2009)
[49] Utgoff, P.E., Berkman, N.C. and Clouse, J.A., "Decision Tree Induction Based on Efficient Tree Restructuring," Machine Learning 29(1), 5-44 (1997).
[50] Sumioka H., Hosoda K. Yoshikawa, Y. and Asada, M., "Acquisition of joint attention through natural interaction utilizing motion cues," Advanced Robotics 21(9), 983-999 (2007).

[51] Smith, J. H., Weisbin, C. R., Elfes, A., Lincoln, W., Mrozinski,J., Shelton, K., Hua,H., and Adumitroaie, V. "Human and Robot Task Allocation: An Operations Management Approach for Work-System Productivity," 38th Annual Meeting of the Western Decision Sciences Institute, Hilton Kauai Beach Resort, Kauai, HI, April 7-11,) 2009).