

# AuRA: Principles and Practice in Review

Ronald C. Arkin and Tucker Balch

College of Computing  
Georgia Institute of Technology  
Atlanta, Georgia 30332

## Abstract

This paper reviews key concepts of the Autonomous Robot Architecture (AuRA). Its structure, strengths, and roots in biology are presented. AuRA is a hybrid deliberative/reactive robotic architecture that has been developed and refined over the past decade. In this article, particular focus is placed on the reactive behavioral component of this hybrid architecture. Various real world robots that have been implemented using this architectural paradigm are discussed, including a case study of a multiagent robotic team that competed and won the 1994 AAAI Mobile Robot Competition.

## 1 Introduction

The Autonomous Robot Architecture (AuRA) was developed in the mid-1980's as a hybrid approach to robotic navigation [6]. Hybridization arises from the presence of two distinct components: a deliberative or hierarchical planner, based on traditional artificial intelligence techniques; and a reactive controller, based upon schema theory [2]. It was the first robot navigational system to be presented in this integrative manner [8, 10].

This article reviews the overall structure of AuRA, describing its evolution and strengths. The biological motivations underlying its design are also surveyed. A discussion of the techniques underlying its use and a review of several robotic implementations is then presented, followed by a detailed discussion of a multi-agent implementation that was a winner in the 1994 AAAI mobile robot competition.

### 1.1 Overall Structure

Schematically, the components of AuRA are depicted in Figure 1. Two major planning and execution components are present: a hierarchical system consisting of a mission planner, spatial reasoner, and plan sequencer, coupled to a reactive system, the schema controller. In the style of a traditional hierarchical planning system as found in the intelligent controls community [1, 38, 47], the highest level of AuRA is a Mission Planner concerned with establishing high level goals for the robot and the constraints within which it must operate. In AuRA-based systems constructed to date, the Mission Planner has acted primarily as an interface to a human commander. The Spatial Reasoner, originally referred to as the Navigator in [6], uses cartographic knowledge stored in long-term memory to construct a sequence of path legs which the robot must execute to complete its mission. In the first implementation of AuRA, this planner used the A\* algorithm to search over a meadow map (hybrid free space/vertex graph) representation [7]. The Plan Sequencer, referred to as the Pilot in earlier work, translates each path leg generated by the Spatial Reasoner into a set of motor behaviors for execution. In the original implementation, the Plan Sequencer was a rudimentary rule-based system. More recently it has been implemented as a finite state sequencer. Finally, the collection of behaviors (schemas), specified and instantiated by the Plan Sequencer, is then sent to the robot for execution. At this point, deliberation ceases, and reactive execution begins.

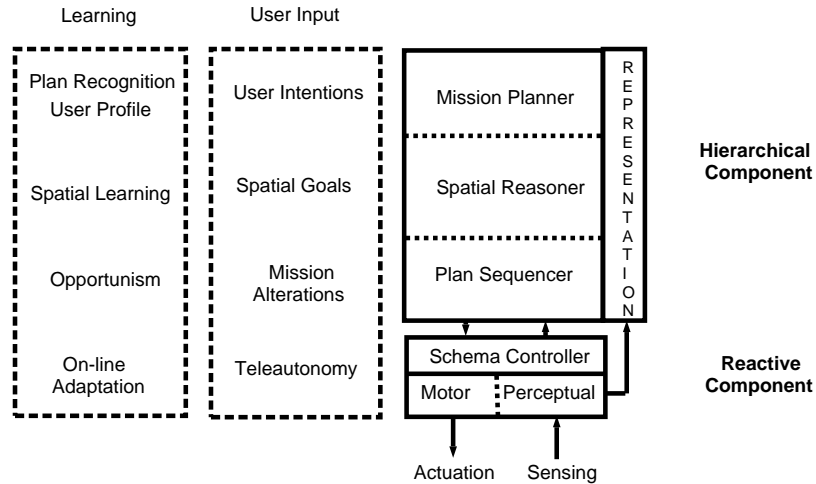


Figure 1: High-level AuRA Schematic

The schema manager is responsible for controlling and monitoring the behavioral processes at run-time. Each motor behavior (or schema) is associated with a perceptual schema capable of providing the stimulus required for that particular behavior. This action-oriented perception is the basis for reactive navigation [11]. Each reactive behavior generates a response vector in a manner analogous to the potential fields method [9]. The schemas can operate asynchronously, transmitting their results to a process (**move-robot**) which sums and normalizes these inputs and transmits them to the low-level control system for execution (Figure 2).

A homeostatic control system [12] (tested only in simulation to date) is interwoven with the motor and perceptual schemas. Internal sensors, such as fuel level and temperature transducers, provide information over a broadcast network which is monitored by behaviors containing suitable receptors. These internal messages change the performance of the overall motor response by altering the relative strengths of the behavior and internal parameters in an effort to maintain balance and system equilibrium (homeostasis).

Once reactive execution begins, the deliberative component is not reactivated unless a failure is detected in the reactive execution of the mission. A typical failure is denoted by lack of progress, evidenced either by a velocity of zero or a time-out. At this point the hierarchical planner is reinvoked one stage at a time, from the bottom up, until the problem is resolved. First, the Plan Sequencer attempts to reroute the robot based on information that has been obtained during navigation and stored in short-term memory. Original implementations used sonar maps produced using the Elfes-Moravec algorithms for spatial world modeling [27]. If for some reason this proves to be unsatisfactory (e.g., the route is completely blocked within this local context), the Spatial Reasoner is reinvoked. It attempts to regenerate a new global route that bypasses the affected region entirely. If this still fails to be satisfactory, the Mission Planner is reinvoked, informing the operator of the difficulty and asking for reformulation or abandonment of the entire mission.

## 1.2 Strengths

Modularity, flexibility, generalizability, and hybridization constitute the principal strengths of the Autonomous Robot Architecture. The value of each of these aspects has been demonstrated in practice both in simulation and on real robotic systems.

AuRA is highly modular by design. Components of the architecture can be replaced with others in a straightforward manner. This is particularly useful in research. Some examples, based on recent or ongoing dissertations, include:

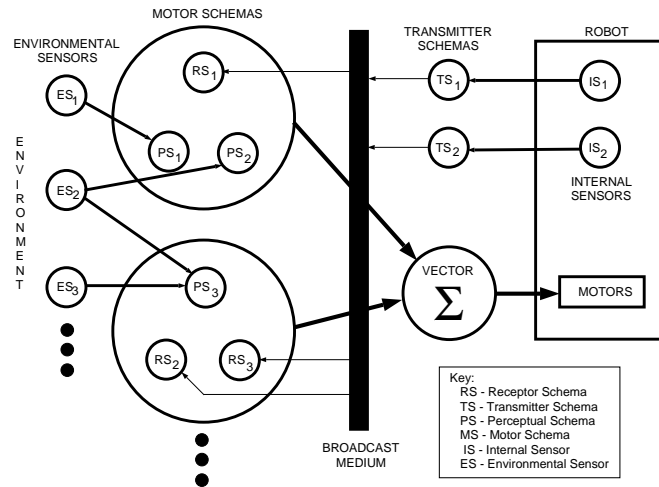


Figure 2: Process diagram for reactive component of AuRA

- A specialized Mission Planner was developed for an assembly task where boxes are pushed together into a specified arrangement. This planner was ported to a Denning mobile robot that competed in the 1993 AAAI Mobile Robot Contest. The planner was further extended in [49] to reason over more general planning tasks.
- The original A\* Spatial Reasoner has been replaced with Router [31], a multi-strategy planner. Router models navigable routes as links between nodes instead of the meadow-map representation used previously. The system was tested on a Denning mobile robot which successfully navigated from room to room and down corridors in our laboratory building.
- Perceptual schemas have been expanded to incorporate specialized action-oriented sensor fusion methods [39]. Recognizing that in many cases multiple sensors sources are better than individual ones, specialized strategies were developed to fuse data together within the context of action-oriented perception. Dempster-Shafer statistical methods provide the basis for evidential reasoning.
- The original rule-based Plan Sequencer has been replaced with a temporal sequencer [16] that traverses a finite state acceptor (FSA) expression of a plan [37]. Each state of the FSA represents a specific combination of behaviors that accomplish one step of the task. Transitions are made from one state to another when significant perceptual events trigger them.

Another strength of AuRA is the flexibility it provides for introducing adaptation and learning methods. In early implementations of AuRA learning arose only from short-term memory of spatial information used for dynamic replanning. Since then, a variety of learning techniques have been introduced including:

- On-line adaptation of motor behaviors using a rule-based methodology [25].
- Case-based reasoning methods to provide discontinuous switching of behaviors based upon the recognition of new situations [43].
- Genetic algorithms that configure the initial control system parameters in an efficient manner [44] and that allow a robot to evolve towards its ecological niche [35] in a given task-environment.

The generalizability of AuRA to a wide range of problems is another strength. Various architectural components have been applied in a variety of domains including:

- Manufacturing environments [18].
- Three Dimensional navigation as found in aerial or undersea domains [14].
- Indoor and outdoor navigation [6].
- Robot competitions [15, 19].
- Vacuuming [36].
- Military scenarios [37, 20].
- Mobile Manipulation [23].
- Multi-robot teams [13, 21].

Several of these robotic systems are described in more detail in Section 4.

Finally, one of the major strengths of AuRA results from the power of wedding two distinct AI paradigms: deliberation and reactivity. The advantages of this strategy have been demonstrated in several other hybrid architectures that have subsequently appeared, most notably Gat’s Atlantis architecture [28], a three layered system.

## 2 Biological Connections

AuRA has, from its inception, been influenced by a wide range of ethological, neuroscientific, and psychological studies. The most profound influence has been that of schema theory [2]. Schema theory is a theory of intelligence which represents motor and perceptual control at a level of abstraction higher than that of neural networks. Developed by Arbib [2], formalized by Lyons [33], and operationalized by Arkin [5], schema theory develops sensorimotor coordination as a series of active concurrent processes, each independently striving to achieve an agent’s goals.

Within AuRA, schemas are employed at the reactive control level, and are encoded using an analog of the potential fields method [32]. Each motor schema receives sensory data from an associated perceptual schema and generates its reaction to the stimulus in the form of a vector. Biological similarities to this form of motor control have been observed in the navigation of a toad [3], and limb control within the spinal cord of a frog [40].

The justification for hybridization of reactive and deliberative control can be found in studies by psychologists such as Norman and Shallice [42] and Neisser [41]. Action-oriented perception also has a strong psychological basis, motivated largely by Gibson’s theory of affordances [30].

The homeostatic control system was developed using models of the mammalian endocrine system as inspiration [12]. These models provide the basis for broadcast communication, the partition of negative feedback control into transmitter and receptor schemas, and the provision for modulation of an ongoing behavioral process by internal sensing.

## 3 Related Robot Architectures

Deliberative hierarchical robotic architectures such as NASREM [1] appeared in the early 1980s. Their continued development in the intelligent controls community has led to a diversity of approaches (e.g., [38, 47]) that are characterized by predictable and predetermined communication between layers and where each layer’s functionality varies in both spatial extent and time criticality. Reactive behavior-based control systems appeared later and emphasized parallelism over hierarchy. Brooks’ subsumption architecture typifies this approach [22] where multiple behaviors are concurrently active and intelligent action emerges through the complex interactions between the behaviors and the world.

Hybrid deliberative/reactive architectures attempted to combine the best of both of these paradigms. AuRA, as described in this paper, first appeared in 1986-87 [4, 6]. Other hybrid architectures followed: including Connell’s SSS [26], Gat’s Atlantis [28], Lyons’ Planner-Reactor [34] and Georgeff’s PRS [29]. Each differs in its overall structure, control methods, emphasis, and interfaces between the deliberative and reactive components. They all, however, share a common philosophy embodying a commitment to integrated yet distinct deliberative planning and reactive control systems.

## 4 AuRA in Practice

We now focus on those systems that have been fielded using the AuRA philosophy. There is a strong emphasis on the reactive aspects of the architecture within our research which is clearly evidenced in the descriptions that follow.

### 4.1 Motor Schemas

Many primitive robotic behaviors have been implemented within AuRA. In particular, some of the motor schemas developed to date include:

- **Move-ahead**: move in a particular compass direction.
- **Move-to-goal** (both ballistic and guarded): move towards a discrete stimulus.
- **Stay-on-path**: move towards the center of a discernible pathway, e.g., a hall or road.
- **Avoid-static-obstacle**: move away from non-threatening obstacles.
- **Dodge**: sidestep approaching ballistic objects.
- **Escape**: Evade intelligent predators.
- **Noise**: move in a random direction for a fixed amount of time. (persistence)
- **Avoid-past**: move away from recently visited areas.
- **Probe**: move towards an open area.
- **Dock**: move in a spiral trajectory towards a particular surface.
- **Teleautonomy** - introduce a human operator at the same level as other behaviors.

Figure 3 shows some of these behaviors represented as fields. The arrow at each point illustrates the vector the behavior would generate if the robot were at that point. The robot does **not** actually compute the entire field, it only computes a vector based on its current local perception. The entire field is shown only to enhance the reader’s understanding. It is important to realize that this approach is far less computationally intensive than traditional potential field methods that require the computation of the entire field. Many of these motor schemas have been generalized to three dimensional navigation [14, 23].

Potential fields are known to suffer from local minima and cyclic problems. The inherent short-sightedness of a purely reactive approach is one motivation for integrating deliberation and reactive control in AuRA. Nonetheless a range of reactive strategies have been developed to circumvent to some degree these inherent low-level problems. These partial solutions include the use of noise as a form of “reactive grease” [9], an **avoid-past** behavior [24], and various methods of learning and adaptation including the use of the notion of momentum [25], case-based reasoning [43] and genetic algorithms [44].

Assemblages of behaviors are formed by combining primitive behaviors in meaningful ways. The relative importance of each behavior is encoded through the use of gain values: scalar multipliers applied to the output of each behavior. The resulting vectors are then combined using vector addition to yield the overall




Figure 3: Example Schema Fields  
 (A) Move-to-goal (guarded) (B) Avoid-static-obstacle  
 (C) Stay-on-path (D) Dodge.




Figure 4: Reactive path generated via schema superpositioning.  
 The active schemas used include stay-on-path, avoid-static-obstacle, move-to-goal, and noise.




Figure 5: Georgia Tech's mobile robots. Left photo, back: George, Ren and Stimpy, front: Ganymede, Io and Callisto.

reaction of the robot to its environment. An example of a path generated via this method appears in Figure 4.

Finite state acceptor (FSA) diagrams are used to represent a sequence of complex behaviors where each state represents an assemblage of schemas. This type of control differs from the reactive systems that preceded in that it does not require arbitration between behaviors and it provides a convenient mechanism for the achievement of multiple concurrent goals. It has no layering of behaviors and consists of a dynamic collection of schemas instantiated for the current context. As the context changes, the planning component can alter the behavioral composition of the robot on the fly if needed. Learning and adaptation are also easily introduced through the flexibility afforded at this level.

## 4.2 Robots

Many robots have been instantiated in the AuRA tradition. They implement varying degrees of the AuRA philosophy, from the lowest level of motor schema control to full integration with deliberative reasoning.

A wide range of sensors have been used within AuRA to provide perception for the robots. These sensors include:

- Ultrasound for obstacle avoidance and object recognition [46].
- Computer vision for goal and/or path recognition.
- Laser bar code reader for localization.
- Infrared proximity sensors for object detection and avoidance behaviors.
- Shaft encoders for position estimation.

HARV was one of the first robots manufactured by Denning Mobile Robotics and delivered to the University of Massachusetts. It is, as are the other Denning Robots, a kinematically holonomic robot which can turn and move in any compass direction. This robot was controlled remotely either via a wire tether or a radio link from a Digital Vax computer. HARV was the testbed upon which the first implementations of motor schemas were investigated [9, 6].

George was the first robot acquired at Georgia Tech when the AuRA project relocated from the University of Massachusetts. It is a Denning DRV-I robot similar to HARV. New motor schemas, including a docking behavior for use manufacturing environments [18], escape, dodge, and an avoid-past behavior for reactive navigation out of box-canyons [24] were initially implemented and tested on George.

Buzz [15] (not pictured) is a Denning MRV-3 that competed in the 1992 AAI Mobile Robot Competition. The task for the competing robots was to explore an office environment and find interesting objects. These objects were free-standing PVC pipes marked with retroreflective tape so that the robots could more easily detect them using computer vision. Buzz uses previously developed motor schemas to seek out and move towards the poles. A novel perceptual schema used visual cues to estimate the location of the poles based on the inherent perspective in the image. Although George was the first AuRA robot to use a temporal sequencing approach for sequencing plans [16], Buzz extended this capability to more complex situations [15].

Ren and Stimpy (the black cylindrical robots in Figure 5) are Denning MRV-2 mobile robots equipped with ultrasonic range detectors and positional shaft encoders. They were recently used to evaluate multiagent communication strategies [19]. FSA encodings of behavioral states for a foraging task were ported to Ren and Stimpy then tested with three different communication modes. The tests validated simulation experiments that measured the quantitative differences in performance between the various types of communication [21]. For these tests the robots were linked to off-board computers, but Ren and Stimpy are now being upgraded with on-board computers for future multiagent research. Ren has also been outfitted with a CRS+ robotic arm for conducting research on mobile manipulation [23].

A similar approach was used in the design of a multi-robot trash-collecting team. The robots, Io, Ganymede and Callisto are visible in the foreground of Figure 5. The forage behavior described above was adapted for the team of three robots. The robots and the FSA for collecting trash are covered in more detail in Section 5.

Multiagent work in AuRA has extended to the development of behaviors for robot formations [20]. In some tasks, such as military scouting, formations help robots minimize sensor overlap and exposure to danger. Formations for two, three, and four robots have been implemented using new motor and perceptual schemas. The behaviors have been tested in simulation, but have also been ported to another architecture (DAMN [45]) for use in Lockheed Martin's Unmanned Ground Vehicles. The formation behaviors are being utilized in two upcoming evaluations of the vehicles: Demo C in the Summer of 1995, and Demo II in 1996.

Most robotic implementations of AuRA have focused on lower levels of the AuRA deliberation and reactivity. Several robots have demonstrated additional higher level reasoning including the integration of various forms of path planning capability. For the 1993 AAI Mobile Robot Contest, Stimpy, a Denning mobile robot, was tasked with pushing boxes into a specified arrangement. The contest rules stipulated that the initial arrangement of boxes and the desired configuration would not be provided a priori, but given only at the time of the contest. For the contest, Stimpy was equipped with a laser bar-code reader for identifying boxes, and a ring of ultrasonic sensors provides for detecting obstacles (the robot is visible on the right in Figure 5) A deliberative reasoner, which forms plans for arranging the boxes was developed and tested in LISP [49]. Each step of the plan was executed by activating an assemblage of motor schemas for achieving it. In case of failure, the reasoner is able to formulate a revised plan and continue. The system ran successfully, but difficulties with sensors led to poor performance on the day of the contest.

Another example of integration between deliberative and reactive components of AuRA concerns the Spatial Reasoner, originally implemented in HARV using an A\* meadow-map algorithm [7]. The meadow-map approach was replaced with Router [31], a multi-strategy planner. Router models navigable routes as links between nodes. The system was implemented on a Denning mobile robot, and tested in our laboratory building at Georgia Tech. The robot, using this integrated system, is able to navigate from room to room and down hallways.

Ren has also been outfitted with a CRS+ robotic arm for conducting research on mobile manipulation [23]. Both the schema based reactive control [23] and the deliberative planning systems were extended to address this more kinematically complex system [17].

Finally the MissionLab simulator [37] has been developed and is available over the world wide web for use at other sites (the URL is <http://www.cc.gatech.edu/ai/robot-lab>). This simulator has an interactive graphic designer and supporting compilers to provide easy generation of schema-based robotic control systems that can be bound to a range of physical robots.



## 5 An AuRA Case Study: Trash-collecting Robots

This section examines one AuRA-based system in depth. In 1994, a group of Georgia Tech students built a team of three robots that search for trash, pick it up, and carry it to wastebaskets. The trash consists of styrofoam coffee cups, wads of paper and soda cans. The problem is made more difficult by real world obstacles like tables and chairs. Our team of robots, named Ganymede Io, and Callisto, won the “Clean Up the Office” event at the 1994 Robot Competition sponsored by the American Association for Artificial Intelligence (AAAI). The competition and other competing robots are described in detail in a special issue of AI Magazine [48, 19]. The three robots are pictured in the foreground of Figure 5.

### 5.1 Trash-collecting Robot Hardware and Sensing

The robots were built using mostly off-the-shelf components in a effort to keep their overall cost low. The power system and computer equipment are enclosed in an aluminum chassis bolted to the top of a treaded mobile base. Primary sensors include bumper switches for collision detection and a color video camera. A custom-built gripper is attached to the front of the robots for grasping trash. An IR sensor mounted in the gripper detects objects close enough for grasping.

Color vision is a key factor in the robots’ success at their task. Since wastebaskets are relatively scarce and do not afford detection by the other onboard sensors, vision is essential for finding them. Vision also allows robots to detect *each other* which is useful in cooperative strategies. A separate primary color was used to identify each important perceptual class: red for trash, blue for wastebaskets, and green for other robots. These colors were chosen because they are easy to find in digitized color images. A blob growing algorithm is used to distinguish noise (small blobs) from actual objects of interest (large blobs). The centroid of a blob is used to localize the detected object relative to the robot.

### 5.2 Low-level Behaviors for Trash-collecting

As in earlier robots implemented using the AuRA philosophy, the lowest level behaviors on Io, Ganymede and Callisto are motor schemas. The Schema Controller instantiates and runs schemas as directed by the Plan Sequencer. At any given instant, a specific group of schemas, composing a behavioral assemblage, is active. Each motor schema computes a vector which indicates a desired direction of motion. The vectors are combined (added), then clipped to generate the overall movement vector sent to the robot’s actuators. Some components of the task, like avoiding collisions, are more important than others. To provide for these differences, a gain value is associated with each schema. The relative values of these gains permit a planner (or human designer) to denote each schema’s relative importance. The vector output of each schema is multiplied by its gain value before the vectors are summed for output to the actuators.

As an example of behavior design in AuRA, consider the **move-to-trash** assemblage, intended to move a robot towards an item of trash while it avoids obstacles. Three perceptual schemas and two motor schemas are instantiated:

- **detect-red-blob**: a perceptual schema that uses vision to find the location of the goal, an item of red trash.
- **detect-obstacles** a perceptual schema that detects and tracks obstacles in the environment using bumper switches.
- **move-to-goal**: a motor schema that generates a vector towards the trash found by **detect-red-blob**.
- **avoid-static-obstacles**: a motor schema that generates a vector away from any detected obstacles (magnitude varies inversely with range to the obstacles).
- **detect-IR-beam-broken**: a perceptual schema that indicates when the IR beam in the robot’s gripper is obstructed. This is used as a trigger to close the gripper around the trash object.




Figure 6: Robot Behavioral State Diagram

The Plan Sequencer will terminate this assemblage and instantiate a new one when **detect-IR-beam-broken** signals an object has entered the robot's gripper, or if the object is lost from view for a period of time. When this transition occurs, all of the schemas are halted and deleted from working memory. Newly instantiated schemas are then copied from a schema library into memory and scheduled to run. Connections between perceptual schemas and motor schemas are initialized as the schemas are loaded. These links permit the use of generic schemas. For instance, when a robot should move towards red trash, the output of a red-detecting perceptual schema is connected to **move-to-goal**, so the robot seeks out red objects. When it should move towards blue objects instead, like a wastebasket, the output of the blue-detecting perceptual schema is passed to **move-to-goal**.

### 5.3 A Plan for Trash-collecting

The clean up task can be broken into smaller subtasks that are executed in order: find some trash, move to it, grasp it, find a wastebasket, move to it, drop the trash. The actual plan is a bit more complicated to allow for opportunism and failure recovery (see Figure 6). For this application the plan was coded by humans. The plan is a sequence of behavioral assemblages and perceptual triggers which cause transitions between them. It is conveniently expressed as a Finite State Acceptor (FSA). The states are identified with circles, and perceptual triggers are directed arcs between them. In each state, a separate behavioral assemblage is active. When the condition indicated on one of the arcs is met, the robot transitions to a new state and behavior. In this AuRA system the FSA is coded using a language called BHDL (Behavior and Hardware Definition Language). The BHDL file is interpreted at run time by the Plan Sequencer. Higher levels of control, like reasoning and planning, were not required for this task.

Initially, the robot begins in the *start* state. No motor schemas are active here, and the robot is waiting

for a signal to begin execution of the rest of the plan. The trigger is a tap to one of the bumper switches. When this happens, the robot starts looking for trash as it executes the behavioral assemblage in the *wander-for-trash* state. At the competition, trash items were plentiful, so a sit-and-spin approach was used; a 360 degree visual search often results in detected trash. Two events may trigger the robot to transition to a new state: trash is visually acquired or the gripper's IR beam is broken. Both of these events indicate that a trash item may have been found.

In the case of visually acquired trash, the robot switches to the *move-to-trash* state (described earlier). In the case of the IR beam being broken, the robot attempts to grasp the trash by closing its gripper. The robot's perceptual processing is not rich enough to directly distinguish between trash items and table legs, so the robot must conduct a short experiment to differentiate between the two. After the gripper closes, the robot backs up. If the item blocking the IR beam is no longer present, we conclude it is an obstacle, otherwise it is trash. In the case that it is trash, the robot executes the *wander-for-trashcan* state, which is similar to the *wander-for-trash* state. In the case that it is not trash, the location is marked as an obstacle in short-term memory and the robot continues its search.

After collecting trash and finding a wastebasket, the robot moves towards it. When vision processing indicates the robot is close to the trashcan, it drops the trash and backs up. Finally, the robot begins its search anew.

## 5.4 Cooperation in Trash-collection

The behaviors for trash-collection described so far make no provision for cooperation. If cooperation is to exist between these robots, it must be implemented without communication, since they are not equipped with communication devices. Previous work at Georgia Tech has shown that cooperation may arise in multiagent foraging teams even without communication [13]. A key aspect of cooperation without communication in the previous work involves recognition of fellow team members in order to spread the team apart during the wander phase. This allows more efficient search of the environment for interesting items like trash. To provide for this, Io, Ganymede and Callisto were painted bright green since that color is readily distinguished in digitized images. The robots were programmed to be repulsed from green objects while in the **wander-for-trash** state.

Io, Ganymede, and Callisto collect trash quite well. They are able to search out and grasp red trash objects; Coke and Sunkist soda cans work best. Other trash that they may stumble upon is automatically grabbed opportunistically. The robots are even able to discriminate between table and chair legs that trip their IR beams and trash: tables resist being carried away, while trash does not. In some tests, the three robots gathered 10 to 15 soda cans and delivered them to blue wastebaskets in under 20 minutes. Inter-robot repulsion in the wander phase is obvious in the robots' motions, but we have not yet quantitatively measured the degree of cooperation this provides.

Some unexpected and humorous robot behavior resulted from our team's effort at tuning the robots for the AAAI competition. After the team arrived at the competition site, they decided to use AAAI provided black wastebaskets, rather than incur a penalty for using the proven blue ones. This change required a careful readjustment of image processing routines. After readjustment, the robots were indeed able to find black wastebaskets, but they also confused dark areas under tables and other robots with wastebaskets. The results included robots apparently offering trash to one another (robot-to-robot handoffs) and trash hidden under tables. Still, the robots managed to win the contest in spite of their unexpected behavior.

## 6 Summary and Conclusions

The Autonomous Robot Architecture provides a framework for the conduct of a wide range of robotic research including deliberative planning, reactive control, homeostasis, action-oriented perception, and machine learning. AuRA has been motivated but not constrained by biological studies, drawing insight wherever available as a guideline for system design.

AuRA's strengths lie in its modularity, permitting ready integration of new approaches to various architectural components; flexibility as evidenced by the ease of introduction of various learning methodologies and novel behaviors; generalizability demonstrated by its applicability of a wide range of domains including robot competitions among others; and most importantly, the use of hybridization to exploit the strengths of both symbolic reasoning and reactive control.

## 7 Acknowledgements

AuRA's robots were developed by a long list of students over the years including: Khaled Ali, Gary Boone, Jonathan Cameron, David Cardoze, Russ Clark, Tom Collins, Harold Forbes, Doug MacKenzie, Robin Murphy, Mark Pearson, Elizabeth Nitz, Juan Carlos Santamaria, Eleni Stroulia, David Vaughn, and Bill Wester among others. The authors are indebted to their research efforts.

Research funding for this work has been provided by ARPA, ONR, NSF, U.S. Army, Westinghouse-Savannah River Technology Center and the Georgia Tech Materials Handling Research Center.

## References

- [1] Albus, J., McCain, H. and Lumia, R., "NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM)", NBS Technical Note 1235, Robot Systems Division, National Bureau of Standards. 1987.
- [2] Arbib, M., "Schema Theory", in *Encyclopedia of Artificial Intelligence*, 2nd ed., (S. Shapiro, ed.) Wiley, 1992, pp. 1427-1443.
- [3] Arbib, M. and House, D. 1987. Depth and Detours: An Essay on Visually Guided Behavior. *Vision, Brain, and Cooperative Computation*, eds. M. Arbib and A. Hanson. Cambridge, Mass.: MIT Press, pp. 139-163.
- [4] Arkin, R.C., "Path Planning for a Vision-based Autonomous Robot", *Proc. SPIE Conference on Mobile Robots*, Cambridge, MA, pp. 240-249, 1986.
- [5] Arkin, R.C., "Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior", *Proceedings of the 1987 IEEE Conference on Robotics and Automation*, Raleigh, NC, pp. 264-271.
- [6] Arkin, R.C., "Towards Cosmopolitan Robots: Intelligent Navigation in Extended Man-made Environments". Ph.D. Thesis, COINS Technical Report 87-80, University of Massachusetts, Department of Computer and Information Science, 1987.
- [7] Arkin, R.C., "Navigational Path Planning for a Vision-based Mobile Robot", *Robotica*, Vol. 7, pp. 49-63, 1989.
- [8] Arkin, R.C., "Towards the Unification of Navigational Planning and Reactive Control", working notes, *AAAI 1989 Spring Symposium on Robot Navigation*, Stanford University, CA., March 1989.
- [9] Arkin, R.C., "Motor Schema-Based Mobile Robot Navigation", *International Journal of Robotics Research*, Vol. 8, No. 4, August 1989, pp. 92-112.
- [10] Arkin, R.C., "Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation", *Robotics and Autonomous Systems*, 6 (1990), pp. 105-122.
- [11] Arkin, R.C., "The Impact of Cybernetics on the Design of a Mobile Robot System: A Case Study", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 20, No. 6, Nov/Dec 1990, pp. 1245-1257.
- [12] Arkin, R.C., "Homeostatic Control for a Mobile Robot: Dynamic Replanning in Hazardous Environments", *Journal of Robotic Systems*, Vol. 9(2), March 1992, pp. 197-214.
- [13] Arkin, R.C., "Cooperation without Communication: Multi-agent Schema-Based Robot Navigation", *Journal of Robotic Systems*, 9(3), pp. 351-364, 1992.
- [14] Arkin, R.C., "Behavior-based Robot Navigation in Extended Domains", *Journal of Adaptive Behavior*, Vol. 1, No. 2, pp. 201-225, 1992.
- [15] Arkin, R.C., Balch, T., Collins, T., Henshaw, A., MacKenzie, D., Nitz, E., Rodriguez, R. and Ward, K., "Buzz: An Instantiation of a Schema-Based Reactive Robotic System", *Proc. International Conference on Intelligent Autonomous Systems: IAS-3*, Pittsburg, PA., Feb. 1993, pp. 418-427.
- [16] Arkin, R.C. and MacKenzie, D., "Temporal Coordination of Perceptual Algorithms for Mobile Robot Navigation", *IEEE Transactions on Robotics and Automation*, Vol. 10, No. 3, June 1994, pp. 276-286.

- [17] Arkin, R.C. and MacKenzie, D., "Planning to Behave: A Hybrid Deliberative/Reactive Control Architecture for Mobile Manipulation", *1994 International Symposium on Robotics and Manufacturing*, Maui, HI, August 1994, pp. 5-12.
- [18] Arkin, R.C. and Murphy, R.A., "Autonomous Navigation in a Manufacturing Environment", *IEEE Transactions on Robotics and Automation*, Vol. 6, No. 4, August 1990, pp. 445-454.
- [19] Balch, T., Boone, G., Collins, T., Forbes, H., MacKenzie, D., and Santamaría, J., "Io, Ganymede and Callisto - a Multiagent Robot Trash-collecting Team. *AI Magazine*, Summer 1995, 16:2.
- [20] Balch, T. and Arkin, R.C., "Motor Schema-based Formation Control for Multiagent Robot Teams", *Proc. 1995 International Conference on Multiagent Systems*, San Francisco, CA, pp. 10-16.
- [21] Balch, T. and Arkin, R.C., "Communication in Reactive Multiagent Robotic Systems", *Autonomous Robots*, Vol. 1, No. 1, pp. 27-52, 1994.
- [22] Brooks, R., "A Robust Layered Control System for a Mobile Robot", *IEEE Journal of Robotics and Automation*, Vol. RA-2, No. 1, 1986 pp. 14-23.
- [23] Cameron, J., MacKenzie, D., Ward, K., Arkin, R.C., Book, W., "Reactive Control for Mobile Manipulation", *Proc. 1993 IEEE International Conference on Robotics and Automation*, Atlanta, GA, May 1993, Vol. 3, pp. 228-235.
- [24] Balch, T.R., Arkin, R.C., "Avoiding the Past: a Simple but Effective Strategy for Reactive Navigation", *Proc. 1993 IEEE International Conference on Robotics and Automation*, Atlanta, GA, May 1993, pp. 678-685.
- [25] Clark, R.J., Arkin, R.C. and Ram, A., "Learning Momentum: On-line Performance Enhancement for Reactive Systems", *Proc. 1992 IEEE International Conference on Robotics and Automation*, Nice, France, May 1992, pp. 111-116.
- [26] Connell, J., "SSS: A Hybrid Architecture Applied to Robot Navigation", *Proc. IEEE International Conference on Robotics and Automation*, Nice, FR, 1992, pp. 2719-2724.
- [27] Elfes, A., "A Sonar-based Mapping and Navigation System", *Proc. IEEE Int. Conf. on Robotics and Auto.*, San Fran., Calif., pp. 1151-1156, 1986.
- [28] Gat, E., "Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots", */jit Proc. AAAI-92*.
- [29] Georgeff, M. and Lansky, A., "Reactive Reasoning and Planning" *Proc. AAAI-87*, pp. 677-682, 1987.
- [30] Gibson, J.J., "The Theory of Affordances", *Perceiving, Acting, and Knowing*, ed. Shaw, R., and Bransford, J., Erlbaum, 1977.
- [31] Goel, A., Ali, K., Donnellan, M., Gomez de Silva Garza, and Callantine, T., "Multistrategy Adaptive Path Planning", *IEEE Expert*, 9(6), December 1994.
- [32] Khatib, O., "Real-time Obstacle Avoidance for Manipulators and Mobile Robots", *Proc. IEEE Int. Conf. Robotics and Automation*, p. 500, St. Louis, 1985.
- [33] Lyons, D. and Arbib, M., "A Formal Model of Computation for Sensory-based Robotics", *IEEE Transactions on Robotics and Automation*, Vol. 6, No. 3, June 1989, pp. 280-293.
- [34] Lyons, D. and Hendriks, A., "Planning for Reactive Robot Behavior" *Proc. 1992 IEEE International Conference on Robotics and Automation*, Nice, FR, pp. 2675-80, 1992.
- [35] MacFarland and Bossert, *Intelligent Behavior in Animals and Robots*, MIT Press, Cambridge, MA., 1993.
- [36] MacKenzie, D., and Balch, T., "Making a Clean Sweep: Behavior Based Vacuuming", *Working Notes of 1993 AAAI Fall Symposium: Instantiating Real-world Agents*. AAAI, 1993, Raleigh, N.C.
- [37] MacKenzie, D., Cameron, J., Arkin, R., "Specification and Execution of Multiagent Missions", to appear *Proc. 1995 Int. Conf. on Intelligent Robotics and Systems IROS '95*, Pittsburg, PA.
- [38] Meystel, A., "Planning in a hierarchical nested controller for autonomous robots", *Proceedings of the 25th Conference on Decision and Control*, Athens, Greece, 1986. IEEE.
- [39] Murphy, R.R., and Arkin, R.C., "SFX: An Architecture for Action-Oriented Sensor Fusion", *Proc. 1992 Int. Conf. on Intelligent Robotics and Systems (IROS)*, Raleigh, N.C., July 1992.
- [40] Mussa-Ivaldi, F., Bizzi, E. and Giszter, S., "Transforming plans into action by tuning passive behavior: A field approximation approach", *Proceedings of the 1991 International Symposium on Intelligent Control*, pages 101-109, Arlington, VA, August 1991. IEEE.
- [41] Neisser, U., *Cognition and Reality: Principles and Implications of Cognitive Psychology*, Freeman, 1976.

- [42] Norman, D. and Shallice, T., "Attention to Action: Willed and Automatic Control of Behavior", from *Consciousness and self-regulation: Advances in research and theory*, Vol. 4, R. Davidson, G. Schwartz, and D. Shapiro, (eds)., Plenum Press, pp. 1-17, 1986.
- [43] Ram, A., Arkin, R.C., Moorman, K., and Clark, R., "Case-based Reactive Navigation: A case-based method for on-line selection and adaptation of reactive control parameters in autonomous robotic systems", Technical Report GIT-CC-92/57, Georgia Tech.
- [44] Ram, A., Arkin, R., Boone, G., and Pearce, M., "Using Genetic Algorithms to Learn Reactive Control Parameters for Autonomous Robotic Navigation", *Journal of Adaptive Behavior*, Vol. 2, No. 3, pp. 277-305, 1994.
- [45] Rosenblatt, J., "DAMN: A Distributed Architecture for Mobile Navigation", *Working Notes AAAI 1995 Spring Symposium on Lessons Learned for Implemented Software Architectures for Physical Agents*, Palo Alto, CA, March 1995.
- [46] Santamaria, J.C. and Arkin, R.C., "Model-based Echolocation of Environmental Objects", *Proc. 1994 IEEE/RSJ Int. Conf. on Intelligent Robotics and Systems (IROS)*, Munich, Germany, Sept. 1994, pp. 64-72.
- [47] Saridis, G. and Valvanis, K., "On the Theory of Intelligent Controls" *Proc. SPIE Conf. on Advances in Intelligent Robotic Systems*, pp. 488-495, 1987.
- [48] Simmons, R., "The 1994 AAAI Mobile Robot Competition" *AI Magazine*, Summer 1995, 16:2.
- [49] Stroulia, E., "Reflective Self-Adaptive Systems", Ph.D. Thesis, College of Computing, Georgia Institute of Technology, 1994.