

Combining Deliberation, Reactivity, and Motivation in the Context of a Behavior-Based Robot Architecture

Alexander Stoytchev

Ronald C. Arkin

Mobile Robot Laboratory
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280 U.S.A.

September 14, 2000

Abstract

This paper describes a mobile robot architecture that addresses three main challenges for robots living in human-inhabited environments. Specifically how to operate in dynamic and unpredictable environment, how to deal with high-level human commands, and how to be engaging and fun for human users. The architecture has proven useful for controlling mobile robots in man-made environments. The results are reported for a fax delivery mission in a normal office environment.

1 Introduction

The focus of robotics continues to shift from stationary robots in a factory workcell environment to mobile service robots operating in human-inhabited environments. These environments are especially challenging for mobile robots because they are highly dynamic and unpredictable in nature. To be successful, mobile robots need to address three main challenges: 1) to adapt quickly to changes in the environment; 2) to understand high level human commands; and 3) to be engaging and fun to use for people.

This paper presents a mobile robot architecture that addresses the three challenges outlined above. The first challenge is addressed by the architecture through the use of a low-level behavior-based reactive controller. This enables the robot to adjust quickly to changes in the environment without human intervention.

The second challenge is addressed by using a deliberative planner to map high-level human commands into a reactive controller. This is in essence a planning-as-configuration approach [3] that solves some of the problems associated with purely reactive

control.

The last challenge is equally difficult to solve. To be engaging, a robot needs to know what people expect from it and what it should do to please them. At the same time, robots need to prioritize between different goals to limit their interaction with people when necessary. This problem will become more acute in the future when the robots themselves have human shape, i.e., humanoids. Not only would people expect these robots to have the physical capabilities of a human, but they would also expect them to have emotional and motivational human capabilities. The approach outlined here uses motivational variables to address this third challenge.

2 Related Work

Pairwise usage of the three system components (deliberative, reactive, and motivational) outlined above have been used in various robot architectures but very few integrate all three. Furthermore, no research, to our knowledge, has studied in depth the interactions between the three separate components and their impact on the overall robot behavior. Most of the existing work on robot motivations has been done in simulation only or limited to face-expression generating robots [5].

Parker used mathematically-modeled motivations to achieve adaptive action selection in the context of the ALLIANCE robot architecture [9]. The goal of her work is to achieve fault tolerant cooperative control of teams of heterogeneous robots; motivations are chosen as the mechanism for achieving this goal. ALLIANCE uses two primary motivations: impatience and acquiescence. Impatience allows the robot to handle situations when other robots fail in performing their task.

Acquiescence, on the other hand, is used in situations when a robot realizes that it cannot complete a task. In both cases the motivations are represented as functions of time and some tuning parameters.

Robot moods were used by Grange et al [7] to solve the “trapped robot” problem for a museum tour guide robot. In this scenario, the robot is often surrounded completely by spectators that block its way so that it is unable to complete the tour. If such a situation is detected (based on sonar data and a case library) the robot will transition to a frustrated mood/temperament and deliver a stern request to the people standing in its way to move away. Surprisingly, this simple strategy worked quite well. Overall the robot displayed five different moods: *happy/busy*, *lonely*, *tired*, *frustrated*, and *confused*. A fuzzy state model was used to implement a gradual transition from one mood/state to another. Motivation was also used to make the interaction between people and the robot more interesting. In all cases the different moods affected only the way the robot spoke using its speech synthesizer.

Arkin et al [4] describe a robotic system based on an ethological model of the praying mantis. This model defines three *internal variables*: fear (associated with predator avoidance), hunger (related to prey acquisition), and sex-drive (mating related). The action selection module takes into account the values of the internal variables as well as the currently visible environmental stimuli and selects the appropriate action. The behavior associated with the motivational variable with the largest value is enabled and if there is an environmental stimulus associated with it the behavior is executed. If this behavior is not enabled, then the next behavior with the highest motivation is evaluated. The variables are modeled as follows: hunger and sex-drive increase linearly with time while fear remains zero until a predator is seen at which point it is set to a predefined fixed value. When the robot has physically contacted prey or a mate, then the hunger or sex-drive values respectively are set back to zero.

The current study extends this work [4] by introducing a deliberative component. Furthermore, it provides a more complex motivational system that allows modeling of habituation to stimuli, and the specification of circadian rhythms. The resulting integrated architecture was tested on a realistic office delivery task.

3 The Architecture

This section presents a hybrid robot architecture that combines three components: *deliberative planning*, *reactive control*, and *motivational drives*. Each of the three components addresses one of the challenges outlined in the introduction. A high level block diagram of the architecture is shown on Figure 1.




Figure 1: High level system architecture.

3.1 Deliberative Subsystem

The deliberative subsystem is used to initially configure and reconfigure the behavioral control system as needed, incorporating the use of high-level knowledge when it is available and reliable. The current approach utilizes existing metric floor plans of buildings, converting them to an intermediate free space representation that can generate way-points for navigational purposes [1]. The intermediate segments are translated into suitable behavioral assemblages for execution on the real robot. A block diagram of this process is shown on Figure 2.

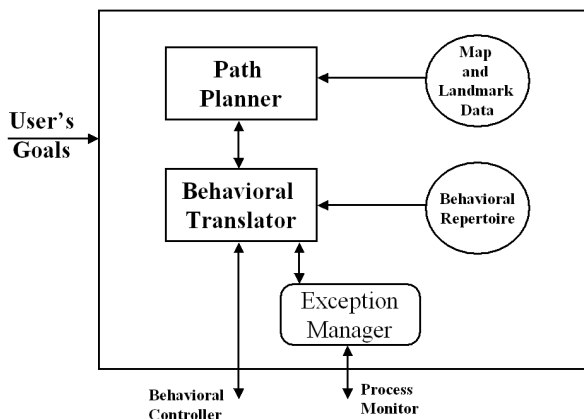


Figure 2: Deliberative Subsystem.

At the core of the *Deliberative Subsystem* is a path planner that takes both user input and an environmental map (in our experiments a map of a building floor) to produce a path from the current robot position to the goal. The map defines the positions of the walls and obstacles in the environment. The method used for path planning is a standard configuration space approach [6]. Once the configuration space is constructed, it is partitioned into a set of non-overlapping convex polygons and the path finding process is recast into a graph search problem. The nodes of the graph are constructed from the mid-points of all convex regions. The A^* search algorithm is then used to find a path between the robot location and the goal point. A path refinement procedure is applied at the end. No claims about the optimality of this approach are made in this paper. However, experience has shown that this approach produces good plans for office environments. One such plan is shown on Figure 3. For more details on this technique refer to [1].




Figure 3: Sample path for an office environment.

The result from the path planner is a sequence of waypoints that, when followed, take the robot to the goal. This sequence of points is used to configure a reactive controller by instantiating a Finite State Acceptor (FSA) using a *Behavior Library* of motor schemas and perceptual triggers used for transitions between states. A diagram of one such controller is shown on Figure 4.

3.2 Reactive control

The behavioral executive is the run-time system that provides the necessary sensorimotor control integration for rapid and intelligent motor response. Schema-based motor control agents [2] that generate a uniform vector representation to encode the responses of the robot to its stimuli. The different agents are easily configurable and can be organized recursively to create more complex behaviors. The depiction of a typical behavioral assemblage is shown in Figure 5. Depending on the complexity of the behavioral coor-




Figure 4: FSA for the path shown on Figure 3. The main building blocks in this case are GoTo motor schemas and AtGoal perceptual triggers.

dinator module, different resulting behaviors can be generated.

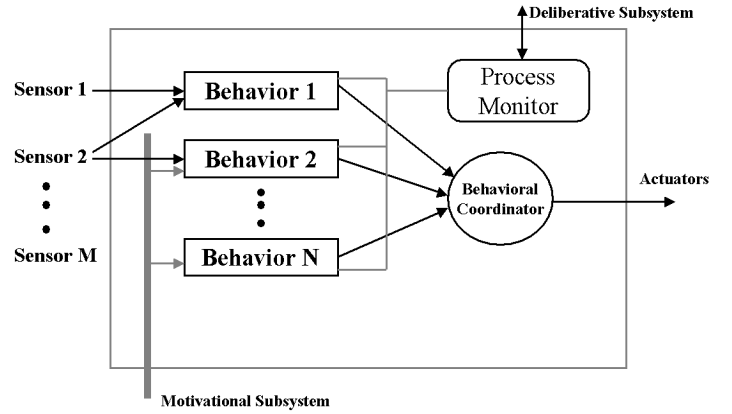


Figure 5: Behavioral Controller: shows the execution of a single Behavioral Assemblage.

The *Process Monitor* monitors the progress of the currently active behavioral assemblage. If no satisfactory progress is being made the current task is aborted and the help of the exception manager in the Deliberative system is requested. The *Exception Manager* can generate a fix to the current strategy. For example, if the robot gets stuck, a localization routine can be invoked after which the original plan will be restarted with a different starting position; i.e., the one found by the localization module.

3.3 Motivational subsystem

The *Motivational subsystem* is responsible for monitoring the internal state of the robot and modulating its behavioral responses as needed through parametric adjustment of the run-time behavioral controller (Figure 6).

3.3.1 Motivational variables

The internal model of the robot is augmented with a set of motivational variables. Real numbers in the range 0 to 1 are used to represent the activation level of each variable. The variables are organized in a motivational vector, but each variable can be read or written independently of the other variables. Both perceptual triggers and behaviors are allowed to access and modify the motivational vector. From an implementation point of view the variables are kept in a database that resides in the shared memory of all behaviors (threads of execution).

3.3.2 Motivated perceptual triggers

Traditional Behavior-based architectures [3] define perceptual triggers that are used to switch from one state of the system to another. The triggering condition usually depends only on the presence of some environmental stimuli that the robot can detect using its own sensors. Less attention has been paid to modeling how the internal state of the robot (motivation) can affect the way it perceives the world.

In this framework, the motivational vector together with the perceptual stimuli are responsible for the firing of the triggers. Furthermore, motivational variables alone can be connected to triggers thus allowing behaviors to be triggered solely by motivations (e.g., hunger and anger).

3.3.3 Motivational Behaviors

Behaviors can also read and write the motivational variables. Reading allows properly encoded behaviors to act differently depending on the values of the motivational variables. Writing allows perceptual input or internal states to modify the motivational variables. Appendix A gives mathematical definitions for these behaviors.

The latter capability makes it possible to model gradual indifference to sensory stimuli. For example, a noise coming from a constant sound source should draw the robot's attention at first, but if the sound continues for a long time it should be ignored and

the robot's attention shifted towards the processing of other more relevant stimuli. This phenomena, called habituation, is widely observed in people and animals [10].

3.3.4 Motivational Processes and Circadian rhythms

Certain motivations change as a function of time in a cyclical manner (e.g., sleepiness). Others change in a temporally dependent manner that is not based on cycles. Hunger, for example, peaks at certain times of the day; while anger, often driven by external events, is likely to decrease over time. The ethological literature describes *circadian rhythms* in great detail [8]. Some rhythms change over very long time periods, e.g., even annually as for migration patterns. In most animals, sex drive changes cyclically as a function of the seasons. Others, like hunger and sleepiness, change on a daily basis.

Within the context of this architecture motivational variables can be altered based on either a circadian rhythm or other time-varying function that can be implemented as a separate thread of execution. Similar primitives used to describe robot behavior can also be used to describe a circadian rhythm. This permits the specification of arbitrarily complex circadian rhythms. Also, multiple circadian rhythms can be specified concurrently.

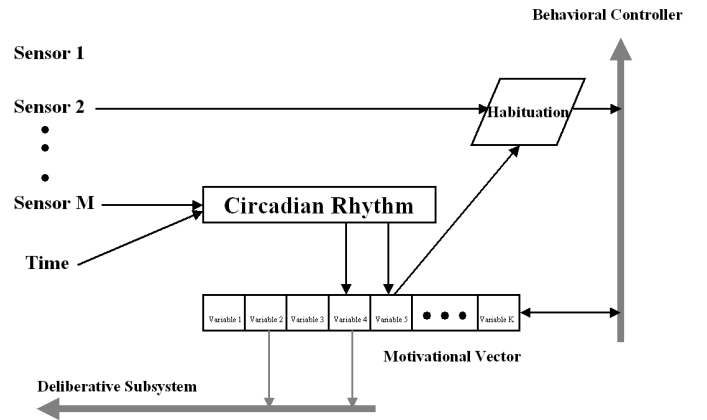


Figure 6: Motivational subsystem.

4 Experiments

In order to estimate the viability of the proposed architecture for expressing useful tasks several exper-

iments were conducted using a Nomad 200 robot, manufactured by Nomadic Technologies, Inc. The robot is equipped with 16 ultrasonic sensors, a Triclops stereo vision system, a Matrox digitizer board, text-to-speech synthesizer, and a binaural sound detection and sound localization system (Fig. 8b). A laptop mounted on top of the robot is used to run the robot control architecture and the GUI interface. Appendix B gives technical details about the software implementation.

Action-oriented perceptual code exists within the system to support obstacle detection using both sonar and a visual bumper (implemented using the Triclops system). Sound detection (direction and magnitude) is available through a DSP board and a binaural microphone pair.

The architecture outlined in the previous sections was used to program a service robot for one floor of an office building. The robot is capable of delivering faxes to users that request the services of the robot, although it is not in regular service. It also can attract people's attention inside the building and ask for assistance if needed. People can also attract the attention of the robot through sound when the robot is moving in the building. The overall task of the robot is to deliver faxes, to interact with people in the building, and to maintain its own internal goals and motivations.

4.1 Human-Robot Interaction

Specifying a task: Using a graphical map interface a human user can request the robot to pick up a fax from the fax room and deliver it to the user. The mapping from the high-level human command to the low-level robot movement commands was explained in Section 3.1.

Getting the Attention of the Robot: While the robot is moving through the building users may try to draw its attention. At present, users can do that by generating a sound (usually through clapping their hands - Fig. 8g). The on-board binaural microphone system is capable of detecting the location and magnitude of the sound relative to the robot's position to within a few degrees. Depending on its internal state, the robot may or may not pay attention to these secondary users (as opposed to primary users that request delivery services).

Getting the Attention of Humans: The robot has a built in speech synthesizer that is used to attract the attention of people standing or walking close to the robot. The robot can utter complete sentences and request favors from the people. For example, when the robot is in the fax room it must request human help

in order to get the fax on-board (the robot has no manipulator currently). Responsive people can help by: getting the fax from the fax machine; putting it on top of the robot; and then pressing a button on the GUI running on the laptop to tell the robot that it now has possession of the fax (Fig. 8e,f).

4.2 Motivational Variables

The robot has four motivational variables: *curiosity*, *frustration*, *homesickness* and *anger*. The variables are associated with the following behaviors. *Curiosity* models the robot's interest in external events (in our case sound events). *Frustration* handles the inability of the robot to complete a task (in our case to get the fax from the fax machine). *Homesickness* makes the robot return to its charging station after it has completed a task successfully or after waiting sufficiently long time before giving up on a current task. *Anger* is linked to behaviors that express the robot's annoyance to external or internal events. Speech behaviors are used to express anger.

Four time-dependent motivational processes are defined and used to gradually change each of the four variables. *Curiosity* and *homesickness* are increased linearly with time, while *frustration* and *anger* are decreased linearly as a function of time.

4.3 Behaviors and Motivational Triggers

While the time-dependent motivational processes tend to bias the motivational variables towards a neutral state, there are several other behaviors that act on these variables in the opposite direction potentially causing motivational triggers to fire and change the current state of the robot.

For example, repeated failure of the robot to achieve its goal of getting the fax triggers a motivational process that increases the value of the *frustration* variable. Sufficiently large values of this variable can cause the robot to stop its current task and make a call to the *exception manager* in the deliberative subsystem (Fig. 2). In the fax delivery scenario, when the robot consistently fails to find a person that can help it get the fax from the fax machine, a new strategy is generated by the *Deliberative Subsystem* which makes the robot leave the fax room and wait in the corridor in front of the room for people to pass by.

Habituation to sound is achieved through the use of the *curiosity* motivational variable. Every time a sound is heard, and attended to, the value of this variable is decreased. Thus, if the sound is continually repeated, the robot stops paying attention to it.

If a user continues to try to attract the attention of the robot after it has stopped paying attention to sound events, the value of the *anger* motivational variable is gradually increased. If this value becomes large enough, the annoying user is warned (using the speech synthesizer) not to pester the robot anymore.

Sometimes the robot may fail to complete a delivery mission because there is no person available to help it. In these situations the robot gives up and returns home. The switch between these two behaviors is guided by the value of the *homesickness* variable which is being updated by its corresponding motivational update process.

4.4 Results

The performance of the entire system was evaluated on an office delivery task as described above. The robot was asked to deliver a fax ten times always starting from the same initial position. This is not a requirement but makes the results from different runs comparable.

The overall mission success rate was used as criteria for success. Out of the ten runs the robot succeeded in achieving its mission in 7 runs. It failed in 3 cases due to accumulation of deadreckoning errors. Moving to more effective localization methods or perceptual event driven navigation (e.g., sign recognition) rather than relying on dead-reckoning alone would solve this problem. The robot found people to help it in 4 cases and in 3 cases it had to abandon its delivery mission and return to its charging station. Figure 8 shows snapshots from a successful fax delivery mission in which people were available to help the robot.

The values of the motivational variables during another mission in which no people were available to help the robot are shown in Figure 7. The jump in the values of *anger* and *frustration* (around T=110 seconds) is caused by the inability of the robot to find a helpful person. This jump causes a transition to the deliberative system which decides that the robot would be better off waiting in the corridor and sends a plan to the reactive component to go there. When even this does not help as there are still no people around, the value of *homesickness* is increased faster than normal (around T=170 seconds) thus making the robot return to its home base. At that point the *homesickness* value is set to zero (T=320 seconds).




Figure 7: Motivational Variables overtime.

5 Summary and Future Work

This paper described a mobile robot architecture that addresses three main challenges for robots living in human-inhabited environments. Specifically, how to operate in dynamic and unpredictable environment, how to deal with high-level human commands, and how to be engaging and fun for human users. The approach taken utilizes a hybrid robot architecture that combines three components: *deliberative planning*, *reactive control*, and *motivational drives*. Circadian rhythms and time-dependent motivational processes are also utilized.

This architecture has proven useful for controlling a mobile service robot in an office environment. Motivational variables were successfully used to affect robot goals and keep the robot focused on its task. Human-robot interaction was also facilitated by the use of motivational variables.

Additional testing remains to be performed to estimate the effectiveness of real world robots that use motivational variables. It seems that people enjoy working with robots that show different moods. However, formal user studies are necessary in order to find out what it is exactly that people like, and how can the robot keep them interested for longer periods of time.

Currently the robot navigates in the environment using a reactive controller but uses only dead reckoning information to keep track of its position. This requires manual initialization of the robot position before each delivery request. Adding an indoor localization module to the system will improve the success rate of the missions and also allow for continuous testing of the architecture.

6 Acknowledgments

This research has been partially supported by Honda R&D, Co., Ltd, a division of Honda Motor Co. Additional funding for the Georgia Tech Mobile Robot Laboratory is provided by DARPA and C.S. Draper Laboratories. The authors would like to thank Ben Wong and Xinyu Li for their help with the implementation of the system.

References

- [1] Arkin, R.C., "Navigational Path Planning for a Vision-based Mobile Robot", *Robotica*, Vol. 7, 1989, pp. 49-63.
- [2] Arkin, R.C., "Motor Schema-Based Mobile Robot Navigation", *International Journal of Robotics Research*, Vol. 8, No. 4, August 1989, pp. 92-112.
- [3] Arkin, R.C., *Behavior-based Robotics*, MIT Press, 1998.
- [4] Ronald C. Arkin, Khaled S. Ali, Alfredo Weitzenfeld, and Francisco Cervantes-Perez, "Behavioral Models of the Praying Mantis as a Basis for Robotic Behavior", *Journal of Robotics and Autonomous Systems*, Vol. 32, No. 1, July 2000, pp. 39-60.
- [5] Breazeal, C., *Sociable Machine: Expressive Social Exchange Between Humans and Robots*, ECE Dept., Ph.D. Dissertation, MIT, May 2000.
- [6] Lozano-Perez, T., "Spatial Planning: A Configuration Space Approach", *IEEE Transactions on Computers* C-32, 2, pp. 108-120, 1983.
- [7] Illah Nourbakhsh, Judith Bobenage, Sebastien Grange, Ron Lutz, Roland Meyer, and Alvaro Soto, "An Affective Mobile Educator with a Full-time Job", *Artificial Intelligence*, Vol. 1-2, No. 114, October 1999, pp.95-124.
- [8] Oatley, K., "Circadian Rhythms in Motivational Systems", in McFarland, D., (Ed.), *Motivational Control Systems Analysis*, Academic PRes, 1974, pp. 437-459.
- [9] Parker, L., "Alliance: An Architecture for Fault Tolerant Multi-robot Cooperation", *IEEE Transactions on Robotics and Automation*, Vol. 14, No. 2, April 1998, pp. 220-40.
- [10] Wang, D., "Habituation", in M. Arbib (Ed.), *Handbook of Brain Theory and Neural Networks*, MIT Press, 1995, pp. 441-444.

Appendix A: Formal Definitions

Let $M = \langle m_1, m_2, \dots, m_k \rangle$ be a motivational vector consisting of k motivational variables. Let S be the domain of all perceivable stimuli and R be the range of robot responses. Then a motivated perceptual trigger can be defined in the following way

$$P = \begin{cases} 1 & \text{for } p(S, M) \geq T; \\ 0 & \text{otherwise.} \end{cases}$$

where T is some threshold value for the perceptual function $p(S, M)$.

Motivated behaviors are defined as a function

$$\beta : S, M \rightarrow R, M$$


Note that in this case M appears both in the arguments and the result of the behavior β expressing the fact that behaviors can modify the motivational vector.

Appendix B: Implementation Details

Three components were available to the *Deliberative subsystem*: two assemblages, *GoTo* and *Localize*, and one perceptual trigger, *AtGoal*. *GoTo* is a behavioral assemblage that takes as parameters the global coordinates of a subgoal point. It consists of three motor schemas: *MoveToGoal*, *AvoidStaticObstacles*, and *Noise* (See [2] for the formal definitions of these behaviors). The *Localize* behavior changes the internal belief of the robot about its location based on user input. The *AtGoal* perceptual trigger fires when the robot is close to its target goal. The closeness is specified by a parameter which was set to 0.5 meters in the experiments.

At the beginning of every experiment all motivational variables were initialized to 0.5. Every 10 seconds these values were updated as follows: *curiosity* +0.045, *frustration* -0.035, *homesickness* +0.04, and *anger* -0.03.


The motivated triggers described in section 4.3 fire based on the values of their corresponding motivational variables. The following thresholds were used: *homesickness* ≥ 0.76 , *frustration* ≥ 0.75 , and *anger* ≥ 0.75 . Sound detection was achieved through a *Detect-SoundMotivated* perceptual trigger which fired only if the sound level was greater than 3 (on a 0 to 5 scale) and *curiosity* was greater than 0.4. *Curiosity* was decreased by 0.015 after each successful sound detection. A *GoToSoundSource* behavior similar to the *GoTo* behavior was used to approach the person generating the sound.



(a)



(b)



(c)



(d)



(e)



(f)



(g)



(h)



(i)

Figure 8: A fax delivery mission: (a) Person 1 requests a fax delivery service from the robot using a GUI interface; (b) the robot moves in the building corridor and (c) enters the fax room; (d) the robot asks person 2 for help with getting the fax ; (e) person 2 picks the fax from the fax machine and (f) puts in on top of the robot; then presses a button on the GUI interface; (g) the robot heads back to deliver the fax but Person 3 wants to interact with the robot and tries to draw its attention by clapping his hands to generate sound; the robot responds by stopping and acknowledging the request for interaction; (h) the robot comes back to the room where it started and announces that there is a fax to be picked up; (i) person 1 picks up the fax.