An Affective Mobile Robot Educator with a Full-time Job

Illah R. Nourbakhsh (1)
Judith Bobenage(2)
Sebastien Grange(3)
Ron Lutz(4)
Roland Meyer(3)
Alvaro Soto(1)

(1) The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
www.ri.cmu.edu

(2) Education Division
Carnegie Museum of Natural History
Pittsburgh, PA 15213
www.clpgh.org/cmnh

(3) Mobot, Inc.
2425 Liberty Ave.
Pittsburgh, PA 15222
www.mobotinc.com

(4) Exhibits Division
Carnegie Museum of Natural History
Pittsburgh, PA 15213
www.clpgh.org/cmnh

Abstract

Sage is a robot that has been installed at the Carnegie Museum of Natural History as a full-time autonomous member of the staff. Its goal is to provide educational content to museum visitors in order to augment their museum experience. This paper discusses all aspects of the related research and development. The functional obstacle avoidance system, which departs from the conventional occupancy grid-based approaches, is described. Sage's topological navigation system, using only color vision and odometric information, is also described. Long-term statistics provide a quantitative measure of performance over a nine month trial period. The process by which Sage's educational content and personality were created and evaluated in collaboration with the museum's Divisions of Education and Exhibits is explained. Finally, the ability of Sage to conduct automatic long-term parameter adjustment is presented.

1 Introduction

Dinosaur Hall (Figure 1) is the most popular exhibit at the Carnegie Museum of Natural History. Next to physically imposing specimens such as *T. Rex* and *Apatasaurus*, there are smaller dinosaurs, such as aquatic reptiles from the Cretaceous Sea, with equally important lessons for the museum visitor. However, few of the smaller exhibits attract their due attention, and so there is a potential for a richer and more complete educational experience.

In the hope of attracting museum visitors to less frequented exhibits, the Carnegie Museum and Carnegie Mellon University's Robotics Institute collaborated on the creation of an educational mobile robot that would serve as a guide and information server in Dinosaur Hall. Robotic mobility presented significant advantages over a stationary kiosk because of the desire to draw visitors to specific exhibit areas. As a tour guide, the robotic educator would attract visitors, then lead them to exhibits of interest, delivering a multimedia presentation at each stop.

Figure 1: Dinosaur Hall at the Carnegie Museum of Natural History

Now, as of nine months after the start of Sage's responsibilities in Dinosaur Hall, this robot has provided 174 days of service to the museum, with 135 of those days consisting of error-free, totally unsupervised operation. Significant reliability challenges have been surmounted to achieve this level of performance. However, this engineering effort does not suffice, for Sage must succeed as an educator; doing so requires solving the social challenge of making Sage a useful member of the staff—a problem rarely encountered by robotics engineers.

This paper discusses Sage in light of both of these challenges. Section 2 begins by describing Sage's hardware, then its underlying obstacle avoidance and navigation software. Long-term statistics are now available regarding navigation and obstacle avoidance performance, and so this section concludes by presenting and evaluating those results. Section 3 describes the process by which Sage's educational content and personality were created in collaboration with the Divisions of Education and Exhibits at the Museum of Natural History. This section then addresses the affective aspect of Sage's behavior, which is intended to facilitate communication between Sage and its viewers. Finally, the section describes the results of a summative evaluation of Sage's efficacy as an educator, as performed by an independent group of researchers from the Carnegie Institute. Section 4 discusses Sage's ability to perform long-term parameter adaptation, research that was only possible once the robot achieved long-term autonomy. Finally, Section 5 offers conclusions and a description of ongoing and future research projects based on this platform.

2 Sage's Prior Competence

The phrase, *prior competence*, was introduced by Mark Drummond to describe the abilities of a low-level reactive system before a higher-level planner is added to the system (Drummond 1989);(Drummond 1993). In the case of Sage, its ability to avoid obstacles and navigate robustly may be seen as a prior competence for Sage's ability to serve as an interactive educator.

Figure 2: Sage

In this Section, Sage's prior competence is described in a bottom-up fashion, beginning with the underlying hardware, then moving up to the obstacle avoidance algorithms and, finally, the navigation algorithms that have been implemented. Because of the availability of Sage performance data over the course of more than half a year of continuous operation, in Section 2.4 we are able to provide an in-depth analysis of the reliability of Sage's hardware, obstacle avoidance algorithms and navigation algorithms.

2.1 Hardware

Shown in Figure 2, Sage is a six foot tall, three hundred pound robot comprised of more than eight major off-the-shelf components. The lower robot housing, extending from the floor to the top of the second sonar sensor ring, is a modified Nomad XR4000 robot (Nomadic Technologies, Inc. Mountain View, CA). The drive system consists of four wheels, each on independent, motorized casters and driven rotationally by independent motors as well. Eight motors drive the four wheels, rotationally and translationally, enabling fully holonomic motion. A common example of this configuration, in a passive form, would be a standard office desk chair, usually with five castered wheels. Because the chair can describe any two-dimensional path and, independently, can rotate around its own axis, it is holonomic.

The XR4000 sensor system consists of rotation and translation encoders on all wheels, touch-sensitive paneling covering the complete cylindrical surface and two rings of sonar sensors, with 24 sonar transducers in each ring.

Computing power consists of a single on-board Pentium P166 running the Linux operating system. This single Pentium processor is responsible for all of the high level functions of the XR4000 (e.g. encoder integration; motor controller communication) and also all of the code developed for Sage (e.g. multimedia control, navigation and obstacle avoidance, vision framegrabbing and processing).

Atop the XR4000 rests the multimedia components and housing, consisting of a Pioneer laserdisc player, an Altec Lansing actively amplified sound system, and an NCR fifteen inch color active matrix display. Finally, protected by the top housing at a height of six feet is a Toshiba KPD50 color CCD camera. A Matrox Meteor framegrabber interfaces this camera to the Pentium processor in the XR4000.

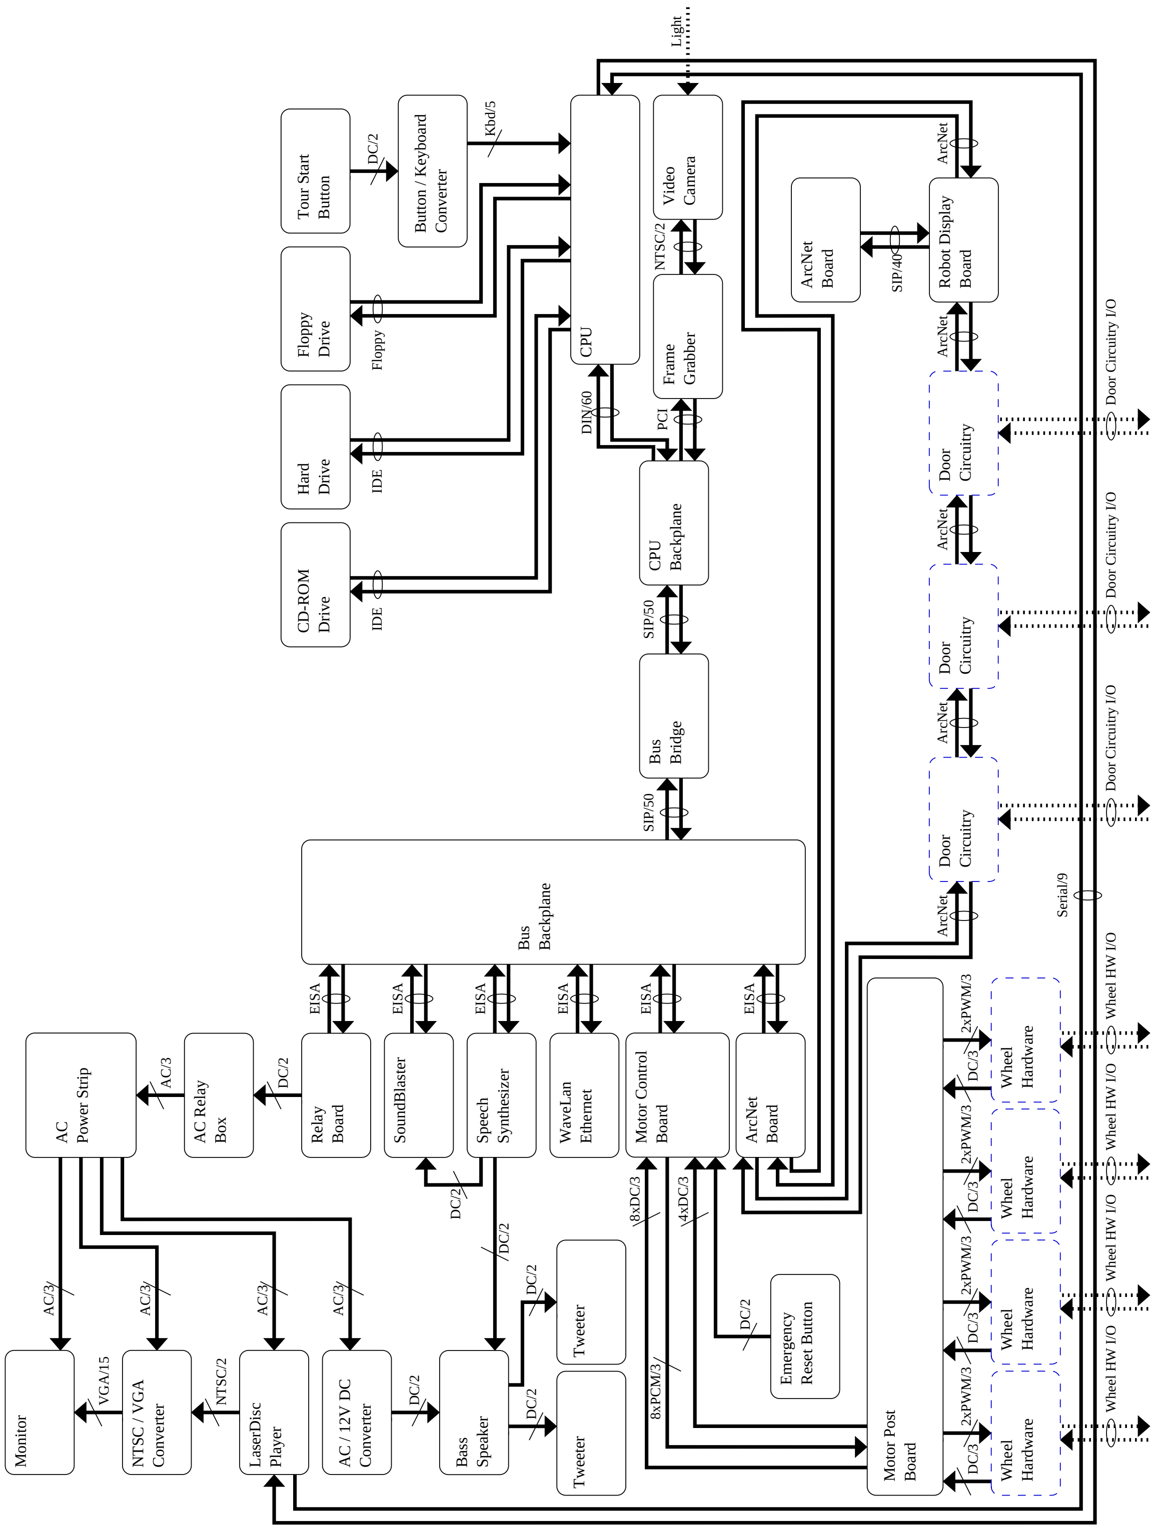

Figure 3: Complete signal schematic of Sage

An emergency stop button is located on the lower surface of the camera housing, protected by a transparent plastic cap. Security guards at the Museum of Natural History are trained in the use of this panic button in case of an emergency. As described below, one of the early consistent causes of robot failure was the use of this button by curious museum visitors. As with every activity that causes the robot to cease functioning, we term the result of such human actions as a robot failure.

The overall power consumption of the entire Sage robot is 300 watts during active operation. On-board battery capacity is designed to comfortably allow for eight hours of activity with 25% charge remaining. A complete schematic of Sage is shown in Figure 3.

2.2 Obstacle Avoidance

In order to enable Sage to achieve the desired level of autonomy, it was an explicit goal to implement an obstacle avoidance system that is sufficiently safe to allow Sage to operate with no supervision. For a three hundred pound robot often surrounded by schoolchildren, such a goal cannot be taken lightly.

In order to enable some degree of code verification, the researchers aimed to create a functional, or completely stateless, obstacle avoidance system with as few instruction lines as possible. Two extremely important physical characteristics of Sage helped to make a stateless and extremely simple obstacle avoidance program attainable:

First, Sage's motion is truly holonomic, enabling changes in the direction of travel that are essentially instantaneous for the purposes of this application. Algorithms that compute trajectories for non-holonomic systems (such as automobiles or electric wheelchairs) can be more complex, since multi-stage maneuvers are often required for entry and egress from confined quarters (Canny 1988; Schwartz & Sharir 1985). In such cases, internal state provides essential meta-level control. For instance, in a three-point parallel parking maneuver, each of three motion planning and control stages is guided by its own subgoal. Taken together, the three motions result in the overall parallel parking goal. Sage has no need for subgoals because of its true holonomic drive system.

Second, Sage has sufficiently complete sensor coverage as to eliminate the robot blind spot entirely. Each sonar ring contains sonar transducers separated by fifteen degrees. Each transducer has a sound half cone of approximately 22.5 degrees, leaving ample overlap between successive sonar pings to locate obstacles that are not perfectly perpendicular to either transducer. The two sonar rings are also positioned along the top and bottom of the XR4000 based, canted down and up, respectively, to provide coverage from approximately 3 inches off the floor to four feet from the floor. In addition, pressure-sensitive microswitches along the entire, vertical outer surface of the XR4000 give Sage a second form of sensor coverage across the robot's lower two-thirds.

In our community, much research has been conducted in creating both transient and permanent environmental maps, such as occupancy grids (Moravec & Elfes 1985; Elfes 1987; Borenstein & Koren 1991), in order to augment limited field-of-view sensors. Vision-based obstacle avoidance systems frequently suffer from a narrow field of view, as do robots with fewer sonar sensors than Sage. MINERVA (Thrun et al. 1999) is a modern tour-guide robot that uses an occupancy grid to compute an obstacle-free direction of travel due to the limited, 180 degree field of view of its SICK laser rangefinder. In such cases, reliable obstacle avoidance may only be possible if the robot constructs a local map, essentially planning its motion as a function of a history of percepts rather than the current, instantaneous percept.

Another, more subtle issue often conflagrates the complexity of occupancy grid-based solutions: once one embarks upon the path of environmental mapping, the differentiation of transient, moving obstacles from sessile, permanent obstacles becomes a necessary skill. Particularly in a dynamic and crowded environment such as the museum environment, approximately 90% of sonar returns are from moving humans who, in the space of less than a second, may be well outside the robot's originally planned path. Failure to recognize these returns leads to overreaction by the robot, which looks random and inexplicable to the untrained human observer. Although many local-mapping obstacle avoidance systems partially address this issue by virtue of short-term memory, we know of no system that can explicitly distinguish, as humans can, between obstacles that should be avoided due to their permanence or instantaneous speed, and obstacles that need not be avoided because they will do much of the avoiding themselves.

Sage sidesteps these challenges by creating no internal representation of the obstacle-ridden environment. Every cycle, Sage computes the most free direction of travel, within a pre-specified window of the desired direction of travel, purely as a function of the current sensor readings. Translational speed is computed independently from the direction of travel, with each computation performed in a single, case-based function with the same control structure as a `cond` statement in Lisp.

```
int calcspeed(int dir, int maxSpeed, int cycleTime)
{
  /* update virtual sonar in the travel direction */
  updateVirtualSonars(dir);

  /* if close obstacle, then stop */
  if (frontBlocked()){
    normal_acc();
    return (0);
  }
  /* else if middle distance obstacle, proportional speed */
  else if (forwardObstacle()){
    middle_acc();
    return (calcSpeedObstacle(maxSpeed-100));
  }
  /* else if NO obstacles for a long distance, full steam ahead !
*/
  else if(clearTowardGoal()){
    smooth_acc();
    return (maxSpeed);
  }
  /* finally, if path clear but far obstacle, reasonable speed */
    else return (maxSpeed - 100);
}
```

Figure 4: C code from the speed computation function in Sage's low-level motion code.

Actual code taken from the forward speed computation demonstrates this approach in Figure 4. The function `frontBlocked()` returns `true` when any forward-facing sonar reads a value of 15 inches or less. Careful inspection of the `frontBlocked()` routine and the code in Figure 4 can help verify that the robot will abide by this most basic level of safety.

The simplicity seen here is representative; the entire obstacle avoidance system consists of just 218 lines of C code.

The functional nature of this code does cause Sage's motion to appear to be extremely reactive to changes in its vicinity. Museum visitors standing in its way will see a smooth and transparently obvious response from Sage when they interact with the robot. But, because of the acceleration parameters set in the low-level motor controller in Sage's drive train, single-cycle, transient sonar returns do not have a noticeable impact on its motion. The obstacle avoidance code will respond, for a single cycle, to such sensor readings; however, the low acceleration values cause a single anomalous cycle to be virtually undetectable. So, a person darting across Sage's path will not cause noticeable changes in its speed or direction of travel.

Of course, motion planning systems such as (Borenstein & Koren 1991; Simmons & Koenig 1995; Burgard et al. 1998) have distinct advantages when unexpected obstacles dictate changes to high-level, planned paths. However, the nature of Sage's venue ensures that such obstacles are nonexistent. Each hallway used by Sage is a public hallway used on a daily basis by museum visitors. Museum staff operate under a policy to keep these hallways clear of obstacles, and therefore virtually every unknown obstacle faced by Sage is either a pedestrian or an infant stroller.

The simplicity of Sage's obstacle avoidance system has been an important factor in its success: over the lifetime of its museum duties, Sage has had zero collisions in a total of 999 hours of running time in motion.

2.3 Navigation

Indoor robot navigation has been a prime research topic in the mobile robotics community for almost three decades. Advances within the last six years have resulted in indoor navigation systems that are far more reliable than their predecessors (Thrun 1995; Castellanos et al. 1997; Murray & Jennings 1997; Simmons & Koenig 1995; Nourbakhsh et al. 1995; Horswill 1993).

Some of the most reliable of these systems have used a form of probabilistic belief update to track the possible positions of a robot over time, thereby allowing for the false ranging values that result from sensor errors, map inaccuracies and unmapped obstacles (Simmons & Koenig 1995; Nourbakhsh et al. 1995; Burgard et al. 1998; Kaelbling et al. 1998; Burgard et al. 1998). The most recent quantitative results in the aforementioned works cite reliability values in the upper ninety percentile range. (Burgard et al. 1998) and (Thrun et al. 1999) describe museum tour guide robots that navigates robustly in the dynamic, uncontrolled environment of a public museum, albeit it under supervision. Such high navigation reliability has been possible both because of more reliable sensing hardware, such as the SICK laser rangefinder, and because of the robustness inherent in probabilistic approaches.

However, such navigation systems are not yet sufficiently reliable for the long-term, *unsupervised* navigation requirements of Sage. For the purposes of this mobile robot educator, our goal was to institute a navigation system with essentially perfect reliability, so that the *prior competence* of the mobile robot at navigation could be taken for granted, and the energy of the research group could be devoted to higher level robot behavior and robot-human interaction issues.

An important second goal of this research was to investigate the reliability of visual navigation. Our hypothesis was that probabilistic navigation is required primarily because of the poor information content of range-finding sensors; visual landmark recognition might obviate the dependence on probabilistic approaches, thereby introducing significant computational and representational savings to mobile robot navigation. The challenge, then, was to create a reliable visual landmark recognition scenario, then test the overall reliability of a purely discrete navigation architecture built upon that visual input alone.

To this end, two concessions were made to surmount Sage's navigation challenge. First, modifications to the environment were allowed, in the form of visual landmarks. Three high-contrast pink rectangles, approximately two feet on a side, were installed on walls terminating three of the hallways traveled by Sage. Using a video camera able to track these landmarks, the robot would be able to navigate those halls reliably. Second, the robot was limited to a predefined set of unidirectional, safe routes, so chosen because those routes are navigable using straight-line paths. Straight paths are followed during the motion control process by means of straightforward, functional algorithms, eliminating the need for more complex motion planning and route following software.

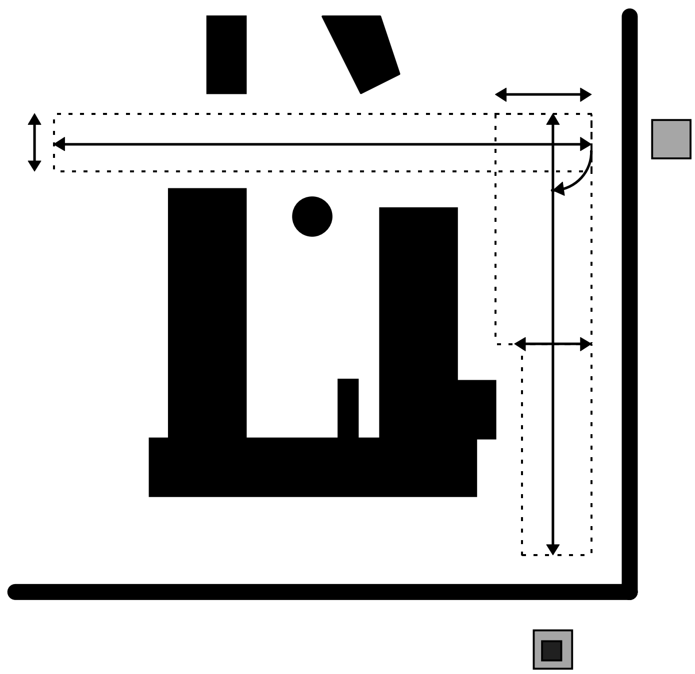

Figure 5: An overhead representation of the information stored in Sage's navigation map for two of its routes. Arrows indicate specific information stored, including the relative angle between routes, the length of each route, the width of safety zones and the type of marker located at the end of each route.

Figure 5 depicts all of the information used by Sage during navigation in Dinosaur Hall. Each route is marked by means of one or more safety zones, a rectangular area that limits the freedom of movement of the robot. This approach has been used before in (Dugan & Nourbakhsh 1993; Knotts et al. 1998; Burgard et al. 1998) and has been shown to enable a mobile robot to avoid static obstacles that it cannot sense, simply by choosing dimensions for the safety zone that exclude those obstacles.

Only two pieces of geometric information are captured by the map: the relative angle of adjacent routes, and the straight-line distance to be traveled between the two ends of a route. The only additional information captured in the map pertains to whether or not there is a visual landmark at the end of each route and, if so, what type of landmark it is.

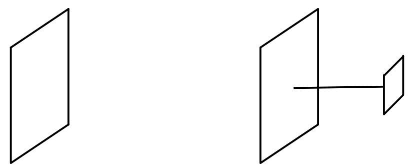

Figure 6: (left) a 2D marker; (right) a 3D marker

Two types of landmarks are used by Sage's navigation system, shown in Figure 6. A 2D marker, as shown in 6a, is simply a flat, rectangular board with a homogeneous, high-saturation color. Such markers are placed at the end of some routes, well above the robot and humans (Figure 7). The exact hue of a 2D marker is not important because Sage is trained to recognize each marker separately. During navigation, the 2D marker conveys simple goal direction information to the robot; by tracking the 2D marker and keeping it in the center of its field of view, Sage is guaranteed to make progress toward the end of the route. Although one could use the precise shape and size of a 2D marker to infer the XY position of a mobile robot operating on a flat surface, Sage undertakes no such computation, instead finding an approximately rectangular region with the appropriate hue and identifying its centroid.

Figure 7: A marker situated at the end of a route from afar (above) and from nearby (below)

The 3D marker, shown in Figure 6b, is born out of observations regarding the three-dimensional markers used by airports to guide airline pilots at gates. When such a 3D marker is placed well above the robot, the relative position of the small black square on the colored background provides the additional information required for the robot to locate its exact position, provided that it can assume purely planar motion. The vertical displacement of the black square provides range data while its horizontal displacement provides lateral position information. Of course, the position of the entire 3D marker in the camera's field of view provides rotational information.

The 3D marker may be seen as a complete local positioning sensor, eliminating any potential for the unchecked accumulation of error over time. The 2D marker provides information regarding direction of travel for individual routes; while this does not ensure global positional accuracy, it is sufficient so long as the robot regularly uses the 3D marker to recalibrate its encoders.

The original intention was to make a single environmental modification, installing one 3D marker to cancel cumulative encoder error. However, empirical testing in the Dinosaur Hall environment showed that the encoders' error accumulation is sufficiently severe to make traversal of more than approximately 10 meters impossible. Further analysis showed that while the encoders can maintain a fairly precise notion of linear distance traveled, they are extremely susceptible to rotational errors. The 2D markers eliminate Sage's reliance on encoders for rotational guidance, and so the robot's strategy involves the used of its camera for rotational guidance while simultaneously using pure encoder data to compute linear distance traveled.

The work of (Lazanas & Latombe 1995) presents a formal framework for this approach, proving that the navigation planning problem can be solved with a sound and complete algorithm when markers exist that provide total local position information. The navigation system used by Sage may be viewed as an instantiation that meets the constraints of this work.

```
While (not (termination())) {
      if (landmarkExists()) {
            /* framegrabber comm */
```

```
                grabImage();
                /* track marker based on last known position */
                trackMarker();
                /* rotation velocity command */
                turnTowardMarker();
        }
        /* obstacle avoidance and goal-directed motion */
        roboMotion();
}
```

Figure 8: Pseudocode depicting the route-following algorithm

Figure 8 shows pseudocode of Sage's basic control loop during route navigation. Because of the robot's holonomic base, the task of rotating to face the 2D marker is independent of the obstacle avoidance code, which aims to travel as much as possible in the direction the robot is facing.

Sage's navigation strategy is consonant with earlier work in which the task of navigating a hallway was seen as a skill to be developed generally and applied to every hallway in the robot's map (Horswill 1993; Gat 1993; Kunz et al. 1999). As with those projects, Sage's navigation reliability depends fundamentally on the reliability of its route-traveling *skill*. The artificial marker that Sage uses makes this route-traveling skill extremely reliable, and therefore the navigation system as a whole is hard to defeat.

This navigation strategy is rather simple; the compromise for this simplicity is a loss of generality. If a landmark is dislodged and falls to the ground, for instance, this robot will fail completely, spending approximately five minutes trying to find the landmark, then giving up, paging us and sending email to us, and shutting itself down. If the robot's path is blocked, and the only available detour involves traveling beyond the edge of the safety zone, it will fail to use that detour.

However, the benefit of this simplicity is reliability. Sage was unexpectedly put to task when a tour group in the museum recognized how Sage was navigating and tried, unsuccessfully, to confuse the robot. What made this tour group especially daunting was that they carried pink flags on poles in order to find one-another. Several pink flags waving in front of its CCD camera certainly gave Sage pause; it actually stopped and waited for the flags to disappear before continuing. Simple temporal filters enable Sage to have strong expectations regarding the position, size and exact hue of the marker being tracked. If the marker found fails to match these expectations, Sage simply waits until it once again does.

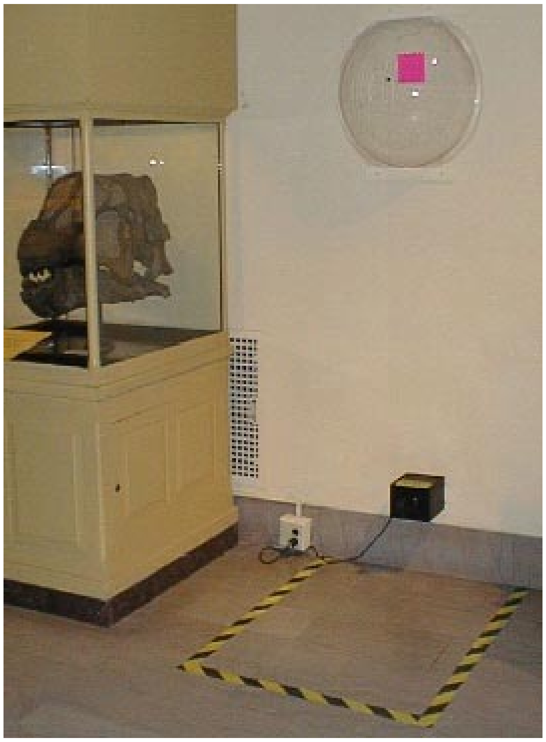

Figure 9: Sage's recharging outlet and, above it, the 3D marker used during docking. In this case, the 3D marker consists of an acrylic hemisphere with a black square painted on it, encasing a flat pink 2D marker.

In order to achieve true self-reliance, Sage also must be able to recharge itself when necessary. This is accomplished using a miniature version of the 3D marker, aligned precisely with a simple wall plug (Figure 9). The original intent was to create a compliant plug so that Sage could dock without a requirement for millimeter-level servoing accuracy during the process. However, once the miniature 3D landmark was built and tested, it was found that Sage could consistently position itself using that landmark with an accuracy of 1.5 millimeters. So, the plug that is now used is a standard, three-pronged computer plug with *no* compliance laterally or vertically. The docking process, which involves moving from over a distance of approximately four meters while tracking and correcting using the 3D marker, takes two to three minutes.

The XR4000 base has the ability to measure the charging and discharging current of the batteries as well as the batteries' independent voltage levels. Using this information, Sage identifies the batteries' state relative to the discharge curve and will begin the charging process when needed. Once docked, Sage also examines the charging current trend to ensure that the charging circuit is indeed operational. If it detects an error, it will attempt several episodes of re-docking and, if the battery charge reaches a dangerous level, it will page us and ask for urgent help (using an alphanumeric pager).

The navigation system described above has met the required goal of zero navigation errors during Sage's operation. However, despite the obstacle avoidance and navigation reliability, Sage's overall reliability in the museum has not been 100%. The next section discloses the sources of Sage's errors over the course of the past nine months, demonstrating the types of problems that occur when a research robot enters long-term, unsupervised operation in the real world.

2.4 Performance

During the period from May 22, 1998 to February 2, 1999 Sage has kept a diary of its performance, including in its records both failures and near-failure anomalies. Table 1 was compiled through the analysis of this eight-month diary. Note that for three of these eight months, covering September through November, the robot was offline due to renovations at Dinosaur Hall.

Table 1: Sage performance resulted for the period 05/22/98 - 02/03/99

Total number of days in the period (save renovation):	174
Total number of totally error-free days	135
Total number of errors	41
Total number of days Sage was down all day	7
Total number of hours Sage is operational	4,008 h
Total number of hours Sage is in motion, giving tours	999
Uptime, last 30 days	97.5%
Mean time between failure, hours	97 (4 days)
MTBF, most recent 30 days (3/15/99 - 4/14/99)	*224 (9 days)*
Mean time to repair, hours	1
Average linear distance traveled by Sage per error free day	1.3 km
Approximate distance covered by Sage for the period	226
Total number of collisions (obstacle avoidance failures)	0
Total number of navigation failures	0

For the purposes of this paper, an error is defined as *any* event that causes Sage to lose its self-reliant ability to provide tours. Therefore, even blatant human mistakes that cause Sage to fail are considered errors. As a matter of both autonomy and convenience, it is important that Sage ask for help whenever a failure does occur. Using an alphanumeric pager and standard email, it will request help in the case of every failure. Indeed, in recent history the only case in which Sage failed without actively requesting our help was an operating system crash on January 11, 1999.

The mean time between failure (MTBF) provides an indication of the frequency with which researchers' attention was required at Dinosaur Hall in order to recover from any Sage error. *As of April 1999, one can expect human intervention with Sage approximately once every nine days.*

Despite the long duration of this project, failures nevertheless continue to occur at a relatively rare rate. Because of the small number of total failures encountered by Sage, it is possible to inspect the complete list of failures, which may lead to insights regarding reliability in long-term autonomous robot projects. Table 2 lists every failure encountered by the mobile robot system. In italics the solution, if applicable, is noted.

Table 2: A comprehensive list of errors. Responses are written in italics.

05/25	Guards turn the lights in Dinosaur Hall off before the robot docks for the night. *Lecture.*
05/28	Robot is herded into and stuck in a concavity. *Narrowed width of the virtual hallway.*
06/02	A 68HC11 sonar controller resets unexpectedly, crashing the robot controller thread.
06/04	Sonar controller resets again. *No solution. Just rebooted.*
06/05	Someone moves the CCD camera's focusing ring, eventually Sage stops. *Lexan cover.*
06/06	Sonar controller resets again. Robot halts mid-tour. *Bug found from yesterday's coding and corrected.*
06/07	Sonar controller resets again. *Adjusted voltage regulator from 5.0 to 5.25 volts and added huge capacitors on the sonar controllers' power lines.*
06/18	Sonar controller resets again. *New thread added that restarts controller process.*
06/20	Sonar initialization fails in mid-tour. *Cause unknown at the time*.
06/21	Relay board (controls power to laserdisc player and monitor) fails. *Remove relay board*.
06/23	Sonar initialization fails. *Bug found: it's the new thread that was added on 6/18. Fixed.*
06/28	Sage is discovered facing completely wrong direction. *Cause unknown.*
06/30	Relay board was reintroduced and fails again. *Changed relay board driver code.*
07/01	Sage faces totally wrong direction again. *Cause unknown. In an attempt to identify cause* *all bells and whistles are turned off or removed.*
07/03	Sage faces totally wrong direction again, but this time prints out garbage, too. *Code pored over. Several small coding bugs found and corrected. One is a memory leak.*
07/09	Sage faces wrong direction. Again. *More poring over of code.* ***Fantastically major malloc() coding error found that explains everything.***
07/18	Sonar initialization fails after sonar controller resets, during error recovery routine. *Bug found in the error recovery routine in the case when the framegrabber also fails.*
07/22	Sonar initialization fails after sonar controller resets, during error recovery routine. *Another bug found in the corrections to the error recovery routine. This time, tested!*
07/22	Break-beam sensor IR transmitter fails. *Replaced.*
07/23	Robot down all day. UPS lost the new break-beam sensor until evening.
07/30	Robot stops and shuts down at the Apatasaurus. *Cause unknown.*
08/01	Robot stops at 3D landmark. ***Voice synthesizer malloc bug found, explains both errors***.
08/07	Noontime lunch break failure: robot didn't fully connect to wall plug. *No action taken*.
08/10	Wall plug fails completely. Batteries drained. *Wall plug worn out. Replaced.*
08/15	Sonar board failure. Reset fails again. *Back to old Nrobot version.*
09/01	Robot calls for help. *Laserdisc player accidentally turned off by me the night before.*
09/08	Renovation of Dinosaur Hall begins. Robot goes on 3 month vacation.
12/02	Robot resumes operation.
12/04	Sudden battery failure causes automatic robot shutdown. *Robot recharged*.
12/05	Same battery fails again, causing shutdown. *Batteries rotated.*
12/10	Robot halted near an obstacle. Cause unknown.
12/11	Laserdisc failure. *Laserdisc power cycled.*
12/12	Motor controller board failure. *No action taken*. Battery problem reappears, charging circuit not functioning correctly.
12/15	Sage is giving tours, but charging rate is questionable. *Much testing ensues*. On 12/16, charging software bug was discovered by Nomadics and corrected.
12/19	Motor controller board failure. *Only hard reboot corrects the problem.*
12/20	Sonar board failure.
12/29	Robot halted near an obstacle. Cause unknown.
12/30	Robot halted near an obstacle. Cause unknown.

01/11	Linux kernel operating system crash.
01/12	Motor controller board failure.
01/15	Sonar board failure. *Problem board identified and replaced.*
01/21	Software update introduced a bug yesterday. *Identified and fixed.*
01/23	Motor controller board failure.
01/25	Robot halted near an obstacle. *Cause of all 4 cases discovered: software bug!*
01/30	Motor controller board failure.

The 41 total failures reported in Table 2 can be classified as 13 unique errors, several of which have recurred. The recurrence has had two causes: in some cases, the source of the error was not identified correctly at first, and so multiple diagnostic cycles were required to correct the error properly. In other cases, the error was identified successfully; however, the correction process introduced a secondary error that then needed to be found and corrected. Statistically, this latter procedure has been the norm and not the exception when software bugs have been discovered on Sage.

A further breakdown of the 13 unique errors indicates that 3 are the results of unexpected human actions (the guard turns the lights off; a visitor defocuses the CCD camera, visitors press the emergency stop). 4 are the result of straightforward software programming errors. Finally, 6 errors are due to stereotypical robotic errors: insufficient agility of the obstacle avoidance module; relay board communication difficulties; failure of the wall plug and the break-beam wheel position sensor; motor controller failures and one operating system crash.

A reassuring trend is that the precipitous drop in the number of new, unique errors over these eight months. Before renovation, a total of 11 unique errors occured. During the two months following renovation, only 2 additional, unique errors have appeared.

Table 2 is equally instructive in identifying what failures have *not* occurred. Despite Sage's simple obstacle avoidance module, which does no global path re-planning, Sage has never encountered unforeseen obstacles in a route and thereby failed. Despite the well-known inaccuracy of robot encoders, Sage's encoders have never performed sufficiently poorly to inhibit Sage's ability to travel a subsequent leg of its tour. Sage has never failed to plug itself successfully into its wall socket for recharging, although the noted failure involved a poor connection between that wall socket and Sage due to a worn plug.

Lessons to be learned may be distilled to two important points: The first is that, whenever a software fix is introduced, we should generally expect that a secondary error will result within the subsequent 24 hours. Only after the second software debugging session can the system be reasonably expected to perform without error.

Second, the hardware failures that occur are generally those failures that are not even imagined beforehand. If our research group had designed a diagnostic engine for Sage, failures would have included sonar transducer failures, stuck-on skin micro-switches, and encoder failures. However, we would never have imagined that the 6811 controlling the sonar transducers would spontaneously reset, causing the motor controller to block; or, that the break-beam wheel position detector's IR emitter would fail (infrared emitters are notoriously reliable). This second lesson learned is somewhat foreboding when one considers automated diagnostic tools for long-term robotic applications. Capturing such “unimaginable” errors by default would have required an arcane representation of the robot’s architecture.

A recent, particularly troublesome error consisted of Sage reporting an inability to move, and eventually timing out and performing a shutdown after signaling its situation to us. This has happened on 1/25, 12/30, 12/29 and 12/10. After many hours of diagnosis, it was discovered that a bug in the original obstacle avoidance code was responsible. The function, `UpdateVirtualSonars()` (Figure 4) contained a bug that would create this error condition when the robot is closer than one meter to its goal position and there are significant numbers of obstacles blocking the robot’s path. This bug dodged discovery over approximately seven months of continuous operation.

This concludes the discussion of Sage’s *prior competence* and its reliability during its first nine months of operation. But this competence only serves as the point of departure for the research projects that make use of Sage. Next, Section 3 addresses the architecture and processes that led to the development of Sage’s core activity as an affective robot educator.

3 The Affective Educator

The social challenge of this project was to create an educator that engages its audience, then effectively provides educational information. The presentations would focus, not only on several of the most popular exhibits (*T. Rex; Apatasaurus*), but also on several smaller and, thus, less frequented exhibits (*Aquatic reptiles*, *Camarasaurus*).

This presents challenges both in terms of educational content development and educator-visitor interaction. As mobile robot designers and programmers, the researchers at The Robotics Institute are certainly not qualified to address these issues alone, and so the educational and affective portions of Sage were designed and developed in collaboration with the Education and Exhibits Divisions at the Carnegie Museum of Natural History.

One of the goals of this article is to describe that collaboration, and thereby demonstrate the fruitful results possible through collaboration between the robotics and the education communities. For this reason, this Section elaborates on the collaborative style used during the design and evaluation cycles of Sage. Section 3.1 describes the creation of Sage’s educational multimedia content, then Section 3.2 details the formative evaluation cycle used on Sage. Section 3.3 describes the creation of an affective architecture for Sage; and, finally, Section 3.4 discusses evaluation techniques used to measure the educational efficacy of the overall system.

3.1 Educational Content

Sage’s original job description called for it to provide an informative 30 minute tour of Dinosaur Hall, taking visitors to a number of exhibits throughout the hall and providing a narrative at each stop. The tour chosen by the Education Division consisted of six tour stops, each with a five minute presentation. Table 3 lists the six stops that were chosen for development.

Table 3: A list of Sage's tour stops and topic sentences.

Tour stop title	Content summary
Introduction	*Provide historical context on the age of dinosaurs.*
Tyrannosaurus rex	*Demonstrate the ferocity and weaknesses of T rex.*
Kansas sea monsters	*Introduce visitors to the Cretaceous Sea.*
Why are dinosaurs so big?	*Present several theories regarding dinosaur size.*
Two-headed *Apatasaurus*	*Summarize the Apatasaurus skull controversy.*
Extinction	*Present and discuss extinction theories.*

In developing a script for each tour stop, the Education Division selected content and created a narrative that would engage and educate viewers across a broad age range. Humor was used effectively, as shown in Figure 10, which contains an excerpt from the script written for the *Kansas Sea Monsters* stop.

> ...Are you wondering how these huge marine reptiles ended up in the American Midwest? I mean, when was the last time someone in Kansas had to swerve to avoid hitting a halibut? Actually, during the Cretaceous period, warm shallow seas covered a wide path of North America from the Arctic Circle to the Gulf of Mexico. Rich in life, the sea was home to huge marine reptiles, ammonites, and prehistoric fish. Some were larger than school buses! The most familiar Cretaceous critters under the sea were plesiosaurs and ichthyosaurs, fully aquatic animals who had some similarities to modern mammals, such as dolphins and whales. Some had large flippers and shark-like tails. Many sported nostrils on the tops of their heads. Think of that swimming around Kansas! The floodplain that bordered the continental sea offered an ideal subtropical climate—sort of a Cretaceous Club Med. Dinosaurs shared their digs with a kaleidoscopic array of present-day creatures such as crocodilians, salamanders, turtles and gars....

Figure 10: A segment of the educational script recorded for Sage.

Next, multimedia presentations that accompany the narrative were produced by the Exhibits Division of the Museum of Natural History. A variety of source material was used, including computer animation, stock footage of paleontologists, re-creations, sound effects and still images.

The script was narrated by a well-known public television personality, and the video and audio were finally assembled and pressed onto a 30-minute optical laserdisc. This technology was chosen because of its extremely high output video quality and random access capability.

Multimedia production is an involved process, both in terms of time spent on research and editing as well as money spent acquiring rights to stock footage and pressing the laserdisc. Months can easily be spent on this process, which costs tens of thousands of dollars. Because of these costs, the multimedia component of Sage can only be revised rarely and with a significant lead time. Therefore the onus is on the evaluation process, described in the next section, to ensure that recommendations for change are considered and backed up by strong evidence.

At the same time, the software design of Sage is faced with the responsibility of being robust and flexible to changes in the educational presentations, in terms of the duration of presentations, the number of tour stops and the tour path.

3.2 The Formative Design Cycle

Formative evaluation is the process by which a work in progress may be evaluated based on observations made before product completion. This activity contrasts with *summative* evaluations, which are conducted at the completion of the project in order to ascertain the product's final performance.

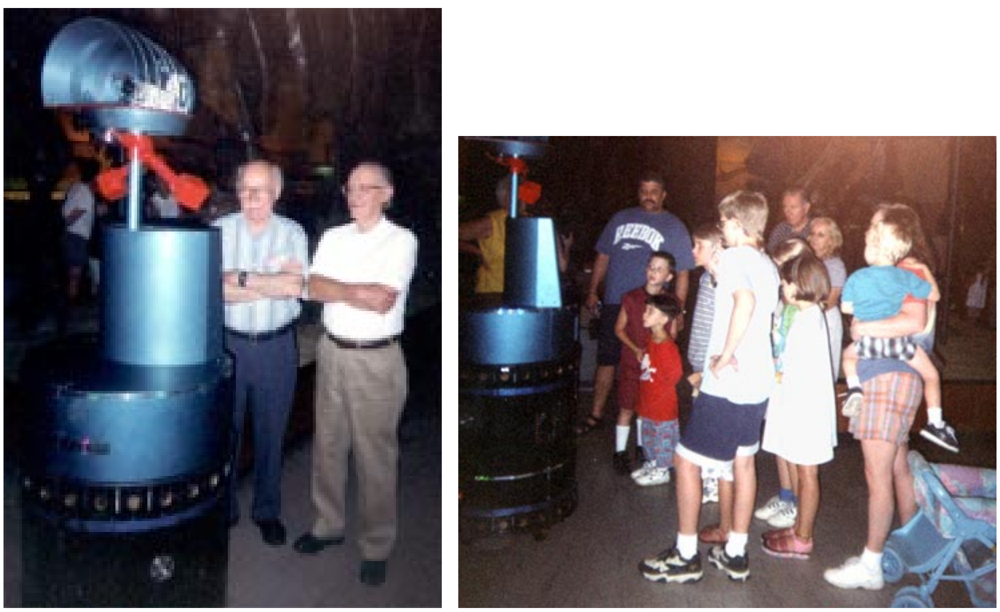

Figure 11: Museum visitors of all ages interact with Sage.

In the case of Sage, a design cycle was instituted, in which formative evaluations by the Education Division served as feedback to Sage's engineers, who then implemented recommended changes to its behavior and personality, and the cycle repeated. The evaluation period consists of two different observation types. One type, *Robot Observation,* involves continuous tracking of the Sage robot during its tours of Dinosaur Hall, recording what is taking place at two minute increments (Figure 11). The second type, *Dinosaur Hall Observation,* identifies and tracks groups throughout all of their activities in the hall.

Once the observations are complete, notes are compiled and conclusions may be drawn regarding the overall impact of Sage on the museum experience. Figures 12 and 13 contain unmodified excerpts of the notes resulting from a Robot Observation session and a

Dinosaur Hall Observation session. Read Figure 12 carefully. The “adult male (CMU?)” individual is an example of some of the dangers awaiting unsupervised, autonomous robots. Happily, Sage appears to have emerged from this brush with a local roboticist unscathed.

2:04	3 adults & 1 child watch monitor
2:06	Same group watching, child follows robot as it moves
2:08	3 children & mother watch. Adult male (C.M.U. student?) touches robot to see if anything happens
2:10	Children lose interest. Adults & kids walk away. CMNH employee & 3 other adults watch robot. Adult male (CMU?) still following.
2:12	Robot moves to next station, family watches closely.
2:14	Adult male (CMU?) gets on the floor to look underneath the robot. 2 children & parents watch and try to guess its name. Adult male (CMU?) drops a pen in front of the robot to see if it is affected by it.
2:16	Robot moves around hall. Adult male (CMU?) follows. Older woman and 2 children follow the robot to the next station. Adult male looks in the “head” of the robot.
2:18	New child approaches the robot, taps the screens & then walks away
2:20	2 children watch robot. Others approach. Adult male (CMU?) still there poking & prodding the robot.
2:22	1 teenage male & 3 younger children watch monitor. Other [sic] walk by curiously but do not stay for long. Adult male still there.
2:24	Father while busily talking on cell phone is drawn to the robot. 2 children watch. Teenage male still watching. Small crowd gathers.
2:26	Robot does not move (has stopped to rewind?). People are still gathered around it.
2:28	Robot begins again. Mother and 3 children, 2 adult males watch, as well as the other adult male (CMU?).
2:30	Two adults watch monitor. Adult male (CMU?) leaves after repeatedly trying to block the path of the robot with his briefcase. He observes how the robot moves around it.

Figure 12: Excerpt from a Robot Observation session.

2:37 Mother (M) and 2 brothers about 4 years (YB) and about 8 years (OB)

“Ah huh. Neat.” From YB as the family entered the hall. They moved between Stegosaurus and Allosaurus with OB pointing to Allosaurus. OB then headed to T. rex and bowed down to it making gestures with his arms. M and YB followed. (No one stopped until they got to T. rex.) Both boys pointed up at T. rex. OB asked M question. “Oh, I see. Come over here.” Said M. Moved under Dippy [Diplodocus] to the T. rex tooth and all touched it. OB spotted Triceratops and said its name. M looked back over the hall pointing toward something in the area of Apatosaurus. They all began walking around the mural and M said, “Look at that guys.” pointing to the turtle skeleton. Then said, “Hey guys, here it is.” about the robot which was behind the mural. (A woman and young girl were at the robot but they left soon after this group moved to the robot.)

2:40 M pointed at the screen, “He’s going to take us. Don’t touch him. It’s going.” They followed the robot. The YB talked but I couldn’t hear what he was saying. The boys were smiling when the robot stopped and turned the corner and started moving. The boys imitated the robot which said “bye, bye.” (As the robot moved along it was joined by 2 older men who left soon after the program began.) The family looked over to T. rex as the program began. They watched silently and the OB looked over at T. rex when the program said, “Check out those teeth.” M also looked up at T. rex then. M and OB laughed out loud at the tooth fairy reference. OB read aloud the please follow message. All followed and looked at Camarasaurus on the way. OB and M looked at Camarasaurus when the program began. M said, “This one here, honey” to YB pointing to Camarasaurus. (YB asked a “how” question about the robot, but I couldn’t hear it.) M put arm around YB as they watched. (Some others, 4 females, passes [sic] by and stopped and YB watched them, not the screen.) OB and M continued watching.

2:46 OB and M smiled at the ‘cave man’ picture. They all followed along with the robot and were joined by another 2 moms and 3 kids. The boys sat at the corner of Apatosaurus exhibit to listen to the video. Boys stood up after a short while. YB looking around hall, looking at other kids, talked to other kids. M and OB watched the screen. M rubbed YB’s head. “Is this one about that one?” YB asked pointing to Apatosaurus. M shook head yes.

Figure 13: Excerpt from a Dinosaur Hall Observation session

The first recommendation tendered by the evaluators concerned Sage’s captioning system. In accordance with the American Disabilities Act, Sage offered close-captioned text to viewers, who would signal their request by pressing an appropriately labeled button next to Sage’s LCD screen.

The Education Division noted in its first Robot Observation that the button was constantly being pushed by children and adults alike, since it was an obvious mode of interaction with the robot; so, it was not being used for the intended purposed by any means. Worse yet, pushing the button so frequently had no noticeable, physical impact on the robot. Even captioning was almost continuously on because the button was a timed ‘on’ switch, not an on-off toggle.

The Education Division’s recommendations were twofold. First, modify Sage so that it always plays the captioned video. Second, modify Sage’s tour (and the button’s label) so that the robot greets visitors beneath the neck of the Apatosaurus and invites visitors to push its button for a tour.

Once the two changes were implemented, further observation revealed that Sage’s interaction with museum visitors improved markedly. First, the tour was given a clear beginning and end from the visitors’ perspective, owing to its additional tour stop under the Apatosaurus neck. Second, visitors receive an immediate visual and audio queue from the robot when they request the tour. As a result, visitors imbued Sage with a greater sense of awareness, and found themselves more engaged and invested in the tour, as they were partially responsible for its launch.

Table 4 provides details of suggestions and modifications that have resulted from this formative evaluation and redesign cycle.

Table 4 –A list of observations and modifications based on the evaluation cycle.

Observation	*Modification*
Tour has no clear beginning	*Sage waits for a request under Apatasaurus neck.*
Captioning button is virtually useless	*Change the role of the button.*
The audience wanders away as presentation spins up	*Begin presentation as robot slows down.*
Confusion about robot's intentions when docking	*Robot tells visitors what it's doing during docking*
Robot continues a tour after the people have left	*Use the sonar detect visitors' arrival and departure.*
The audience attention span is too short	*Decrease tour stop presentation durations.*
The speech synthesizer is difficult to understand	*Replace the synthesizer with digitized sounds*
Some visitors 'abuse' the robot by playing with it	*The robot expresses its awareness of their actions.*

The last item in Table 4 requires elaboration. In one rather common mode of interaction, museum visitors were often observed standing in the robot's path, attempting to either disorient the robot or frustrate its planned direction of travel. In some cases, individuals would even jump up to cover the CCD camera with their hands. This behavior is predictable when directed against a robot; however, it is inappropriate when directed against an educator. Obviously, visitors perceived the robot simply as a robot, not recognizing its educational aspect as a primary identity.

In order to counteract this tendency to treat Sage inhumanely, a decision was made to change the robot's basic architecture to make it an *affective* robot. The idea was that a robot that expresses its goals in an emotive manner may garner greater respect from museum visitors. For instance, a robot that expresses frustration when its educational goals are blocked by pranks may be able to convince those pranksters to treat it as an educator and allow it to fulfill its mission. The goal was not to humanize the robot in order to confuse the distinction between machine and man; rather, the goal was to give Sage the affective tools to communicate successfully with humans, so that it would operate in an atmosphere more conducive to its educational mission.

3.3 The Creation of an Affective Robot

The observations of the Education Division offer clues as to how people interact with a robot in a public place. It was clear that the robot's behavior may impact this relationship between humans and machine. The fascinating question that emerged was: How can one engineer the behavior of a machine in order to establish a new mode of interaction between the untrained public and an autonomous robot?

An affective robot personality was hypothesized as the answer to this question. Establishing this personality would be a team effort of the engineers and the educator/evaluators. First, the engineers created a general architecture that served as a mood transition schema. A fuzzy state machine model was chosen, whereby external events would gradually trigger transitions from one state, or mood, to another. The particular implementation fits in the class of ongoing research on the computational modeling of emotion (Ushida et al. 1998; Velasques 1998; Hayes-Roth & van Gent 1997). Similar to (Ushida et al. 1998), where emotional intensities are computed based on sequences of events, the transition model scores the activation level of each emotion based on events perceived by the robot. As with the Cathexis model (Velasquez 1997), emotion is used primarily to advise the robot's action selection. An improvement over some models involves the blending of multiple emotions, as described in (Velasquez 1998). This blending approach, which leads to smoother transitions in the robot's behavior, is also implemented on Sage.

Once the architectural effort was complete, the Education Division was briefed on the mood mechanism and on Sage's ability to use its sensors to detect events in the Hall. Then, the engineers posed three questions to the educators:

1. What set of moods should Sage have?
2. What events should cause Sage to change its mood?
3. How should Sage manifest its emotional state?

The moods chosen by the Education Division were *happy/busy, lonely, tired, frustrated* and *confused*. The educators also presented scenarios that should cause transitions between moods. For instance, in the case of transitions between *happy* and *frustrated*, one scenario involved the robot's behavior when, mid-tour and with several visitors in tow, another visitor blocks its path. If the person persists in blocking Sage's path, it should transition to a frustrated temperament, at which point it might increase its voice level and pitch and deliver stern requests for the person to stand down: "*I am giving a tour to these visitors right now. Please let me continue!*" In contrast, when a lonely Sage is confronted by a visitor blocking its path, it will be both playful and enticing, engaging the visitor and inviting the person on a tour of Dinosaur Hall.

Sage's affective architecture has been operational as of July 1, 1998. Before this time, it was frequently observed that Sage would be blocked by visitors of all ages, individuals who were essentially testing their ability to influence Sage's behavior. Such blockage would last quite long, sometimes even triggering visitors following Sage to ask the person to step aside. This form of interaction has been noticeably reduced, both in frequency and duration. In particular, when an individual blocks Sage, they are quickly made aware of its intentions as it requests them to step aside. Rather than finding this level of awareness surprising, visitors often accept Sage's requests in stride.

Further evaluations have indicated that the modifications made to Sage have been at least partially successful; visitors appear to be paying greater attention to the educational films and, in some cases, visitors have clearly made the visual connection between the educational material on-screen and the details in the hall that are being discussed. However, the current effort to produce new educational content in the form a greater number of shorter presentations will radically alter Sage's behavior. This milestone will serve as the beginning point for the next evaluation cycle.

3.4 Measuring Sage's Educational Efficacy

In late Fall 1998, individuals from outside of this project conducted an evaluation of Sage's effectiveness (Roehrig and Stockdale 1998). The formal focus of this project was to answer the following question:

Is Chips [Sage's museum name] an effective vehicle to educate visitors in Dinosaur Hall when compared with a docent? Effectiveness is defined as being accessible, educational, entertaining, and appealing to a broad range of visitors.

The Sage project has no intention of replacing docents with robotic tour guides. In our view and in the view of the docents at the Carnegie Museum, their presentations and interactions with museum visitors are complementary. However, for the purposes of measuring educational efficacy, the evaluators required a baseline, and the docent tour provided such quantifiable statistics.

The evaluators chose two methods for collecting data: Robot Observation studies and Questionnaires. Results of the observation studies and questionnaire forms were analyzed with respect to the four effectiveness objectives identified by the team: accessibility, educational efficacy, entertainment and appeal to a broad audience range.

Quantitative measurements of accessibility provide evidence for a general conclusion, that visitors will tend to stay with Sage for a shorter total duration but may return later, whereas visitors tend to follow a docent for the entire tour loop. The team found that 40% of visitors remained with the docent tour for 30 minutes or more, whereas only 4% of visitors remained with Sage for the same duration. However, 74% of visitors remained with Sage for between 5 and 15 minutes, compared with 24% in the case of docent tours. The most significant differences between Sage and docents based on questionnaire results involved tour speed and sound level. Sage's overall speed was viewed more favorably while the docents were rated as easier to hear. Interestingly, these results held for questionnaires returned from both adults and children.

Educational efficacy was measured by asking adults and children knowledge-testing questions both before and after tours. The questions for adults and the success rates before and after robot tours are shown in Table 5.

Table 5: Educational concept questions: success rates before and after robot tours

Question	**<before>**	**<after>**
All dinosaurs lived during the same time period	50%	92%
All dinosaurs were huge animals	50%	72%
Other animals lived on the Earth with dinosaurs	50%	76%
All dinosaurs were carnivorous	48%	80%
All scientists agree on how to put dinosaur bones together	40%	76%
All bones in Dinosaur Hall are real	36%	52%

In the case of the robot's entertainment value, results from both observations and questionnaires rate Sage slightly inferior to the docents. For instance, every visitor rated their enjoyment of docent tours as 'A lot,' and would recommend the tour to a friend with a rating of 'Definitely.' In the case of Sage, the rating was evenly split between 'A lot' and 'Pretty good,' with the recommending Sage to a friend scoring an even split of 'Definitely' and 'Probably,' with a single lukewarm 'Maybe.'

The appeal of Sage to a broad range of audiences was measured primarily by collecting demographic details on museum visitors. On a hopeful note, 20% of Sage's visitors were minorities, whereas the museum-visiting population in Pittsburgh consists of roughly 11% minorities. Peak age groups represented were 5-12 year olds and 25-34 year olds for Sage, while the peak age group for docent tours was 35-44 years of age. Gender differences were also measured, and the results were interesting. Although boys outnumber girls over all tour types, almost by a ratio of 2:1, in the case of Sage, girls were approximately 20% above average and boys were 10% below their average, closing the gender gap somewhat on robot-guided tours.

The results summarized above include a very small portion of the outside evaluation. For the complete, detailed results, refer to the original study (Roehrig and Stockdale 1998).

4 Long-term Parameter Adaptation

A long-term, autonomous robot such as Sage is a prerequisite technology for life-long learning research. In early July 1998, once Sage's autonomy and reliability were established, this project began such a research project in earnest.

As a first step, the goal was set forth to identify performance elements of the robot that would benefit measurably from long-term parameter adjustment. Parameter search is a simple learning process that has the potential for significant gains in performance (Moore et al. 1998), and so it serves as a high payoff first step in lieu of addressing \ more general life-long learning problems.

In the case of Sage, a large number of preset parameters are direct attempts to estimate the values of external world features. Such values can be easy to view as parameter search problems because discrepancies between the external world feature values and Sage's internal parameter settings can often be measured directly using Sage's sensors. Furthermore, many of Sage's parameters are wholly independent, applicable to only one navigable route, and so the computational dangers of joint multiple parameter search can be avoided.

Section 4.1 describes Sage's general framework for defining and evaluating adjustable parameters. Then, Section 4.2 lists the exact parameters on Sage that were expressed as search problems, known as *soft constants*, then adjusted over time. Over the course of several weeks, as Section 4.2 will show, quantifiable improvements in Sage's performance were noted, due to the parameter adjustment module.

4.1 Soft Constants

A soft constant is a limited type of automatically adjustable parameter. The following basic requirements must be met for a candidate in the robot code to be converted to a soft constant:

- The candidate must be treated as a constant in the code.

- There must be some fixed, optimal value for the candidate.
- The transient optimum should be estimable based on the current value and the robot's measured performance.
- A simple mathematical function (e.g. minimum; maximum; mean) should, when applied to a set of transient, estimated optima, approximate the fixed optimal value.

For instance, Sage's navigation map contains information on the relative angles of adjacent routes. These relative angles are treated as constants, and are used to control the robot's rotation when transitioning from one route to the next route. A fixed optimum for the angle exists, and is determined by the ground truth of the actual angles between those routes, compensating for systematic encoder drift errors on the robot.

Each time Sage turns a corner, it locates a new landmark and can thereby estimate the discrepancy between its turning angle and the theoretically correct turning angle, based on the offset of the landmark from the center of its field of view. Therefore, transient measurements of the optimal turning angle can also be computed. Based on these qualities, route relative angles would qualify as good candidates for soft constant instantiation.

In order to transform a constant into a soft constant, the user must provide the following data and routines:

1. an initial value and allowable range for the soft constant
2. a function `estimate` for estimating the transient optimum value based on recent robot performance
3. memory store length specification *l*
4. an update function `update`, specifying minimum, maximum or mean
5. a deviation filter value *c*

Figure 14 depicts the flow of information throughout the parameter adjustment process. The confidence filter is initialized wide open, allowing all estimates to be appended to memory. Once the memory is saturated (i.e. filled with *l* estimates), then the soft constant's value is updated using the *update* function. For instance, in the case of a mean update function, the values of the *l* estimates are simply averaged, and the soft constant's value is set to this result.

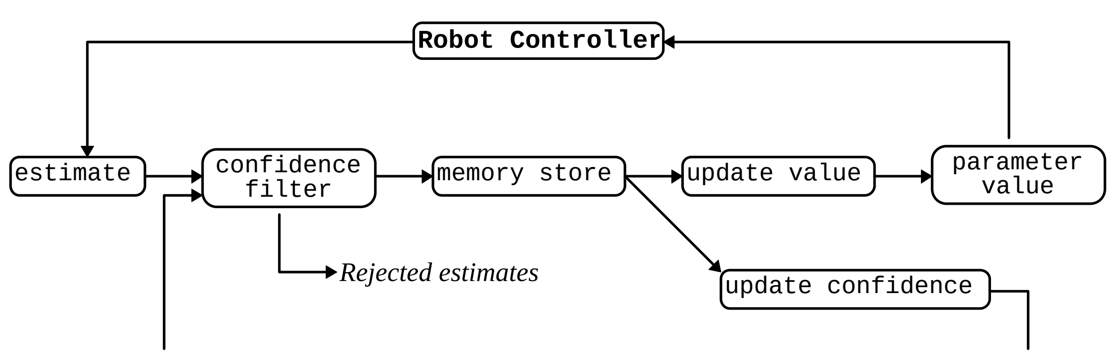

Figure 14: The adjustment process makes changes to the parameter value after every measurement cycle, using a simple hill-climbing strategy as defined by the update function. In addition, the confidence measure is updated every cycle, then used to filter out anomalous estimates from the incoming stream. If this anomaly detection and rejection frequency increases to a user-specified threshold, an error is reported by the soft constant mechanism.

Modifications to the parameter value cause changes to the robot's behavior via the Robot Controller, and thus the feedback loop is closed for parameter adjustment.

An additional, second loop is designed to filter out anomalous measurements. Once the memory store is saturated, it simply operates as a FIFO queue, retaining only the last *l* accepted measurements. At each cycle, the confidence filter is narrowed based on the standard deviation of those *l* estimates and the user-specified value,*c*.

Over time, if the parameter adjustment process increases measurement stability, the confidence filter tends to narrow correspondingly, more stringently rejecting anomalous readings that occur, for instance, when a particularly troublesome teenager keeps Sage busy for several minutes. In such a high traffic situation, the robots travels in sufficient lateral motions to introduce significant rotational encoder error, making the error measurement at the end of the hall unrepresentative of the robot's general error accumulation.

4.2 Parameter Search Performance

The soft constant architecture was implemented for approximately 60 constants in the Sage code. Many navigation routes contain similar sets of constants; factoring in this redundancy leads to a count of 13 unique types of soft constants. Table 6 lists these types of soft constants along with the update function chosen for each.

Table 6: Sage parameters implemented as soft constants and their associated update functions.

Soft constant type	*Update function*
Relative angles of adjacent routes	*mean*
Landmark position for first acquire (*x1, y1, x2, y2*)	*mean*
Left and top edges of search window for acquisition	*minimum*
Right and bottom edges of search window	*maximum*
Search window expansion for marker tracking (*top, bottom, left, right*)	*maximum.*

Beginning on July 20, 1998, the 60 soft constants were introduced in a staggered fashion, with parameter adjustment experiments extending through September 2. Figures 15 and

16 depict the measured improvements of two forms of parameter adjustment: one for image acquisition at the beginning of a navigation route, the other for image tracking during motion through a route. As shown in Figure 15, the miss rate for image acquisition decreased considerably over the test period. This resulted primarily from more precise hallway angles and more accurate position expectations for the markers during the acquisition process.

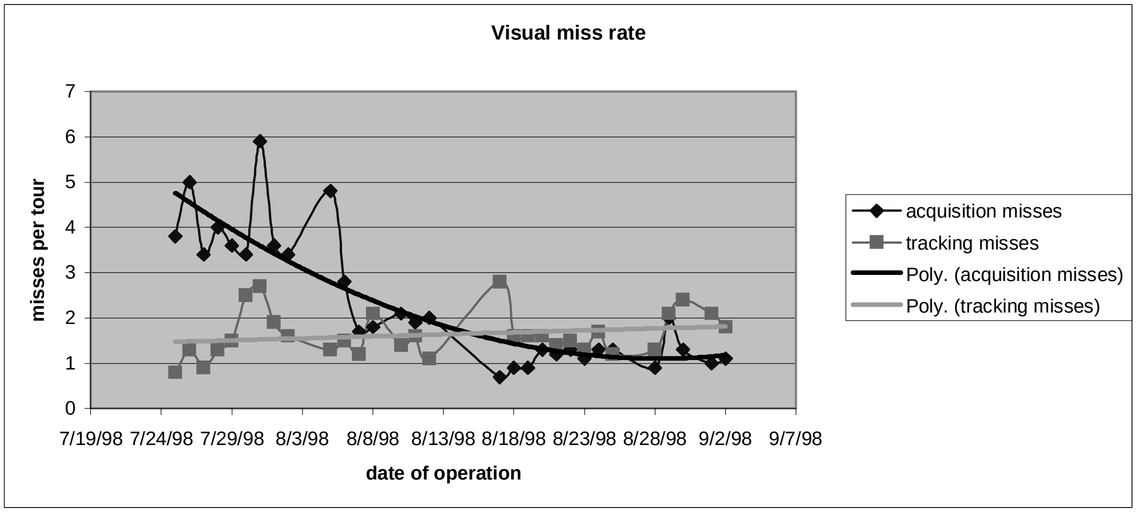

Figure 15: The improving miss rates for static acquisitions and marker tracking over the course of 6 weeks of automatic parameter adaptation.

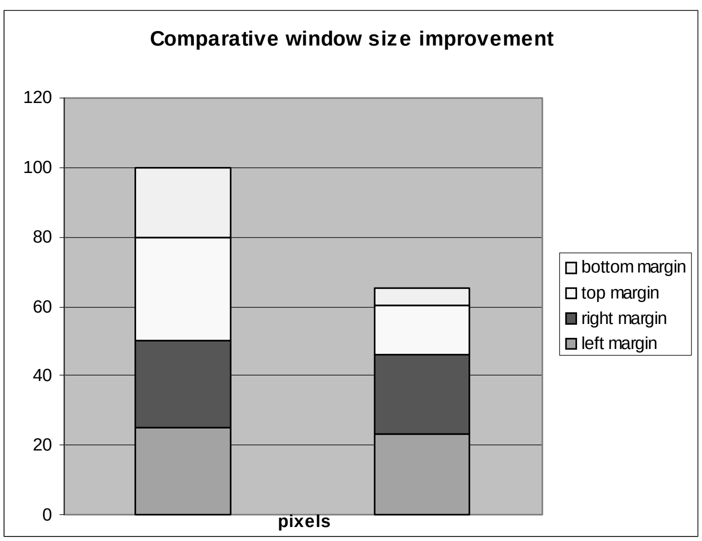

Figure 16: Comparison of the tracking window size before (left) and after (right) 6 weeks of adaptation.

Note that the miss rate for image tracking (Figure 15) actually increased somewhat over the same period. The reason for this trend can be observed in Figure 16, which shows that the window size being used during tracking has shrunk over the same period. This trend was to be expected, as *c* was set sufficiently small to allow the tracking window to decrease in size rather than become inflated to capture the rare large changes in the tracked marker position resulting from large lateral motions of the robot due to obstacle avoidance. Thus the compromise lies between the safety of a small window for tracking, to eliminate the chance for incorrect tracking due to external influences, and the number of tracking misses, which demand a broad and slow search of a much larger portion of the visual field to relocate the marker. This setting is directly related to the length of the memory store.

The quantitative improvements gradually made over the course of six weeks of adaptation is pleasing. However, the soft constants architecture is perhaps the simplest form of parameter adjustment possible. In future steps, we intend to apply search techniques for multiple dependent parameters to Sage. One robot competency that stands to gain significantly in this regard is Sage's docking procedure, which currently requires approximately three minutes for three meters of visual servoing. With careful tuning of the velocity and timing parameters, it should be possible to servo into the wall plug in one smooth motion, reducing the docking time by as much as 150 seconds.

5 Conclusions

Throughout this paper there have been references to RHINO and MINERVA, two other museum tour-guide robots (Burgard et al. 1998; Thrun et al. 1999). The recent introduction of three such robots, all in the same venue, begs comparison. The first distinction that should be drawn involves the research goals of each project. Both RHINO and MINERVA use a probabilistic approach to navigation. These applications demonstrate that the methods used are extremely reliable in spite of the dynamic nature of a public museum space. The Sage project is fundamentally an attempt to create a truly lifelong robot, in answer to the challenges posed by Nils Nilsson in (Nilsson 1996). Our success criteria is a robot that operates without human intervention over the *long term*, where our goal is expressed in terms of months and years of autonomous interaction with the public.

This difference in goals gives rise to a bias in our work towards *simplicity* in all aspects of the architecture. This bias is evident in many aspects of Sage: its single-threaded control code; the case-based obstacle avoidance strategy; a topological world representation; et cetera.

The most unconventional manifestation of this bias is Sage's use of artificial visual landmarks to enable a discrete navigation system employing color vision alone. This approach stands in contrast to MINERVA and RHINO, both of which navigated with no special modifications to their venues. Our conjecture is that the computationally expensive process of probabilistic navigation may be obviated by vision-based sensing. The Sage robot lends credence to this, providing an existence proof that perfectly reliable navigation is possible using vision alone, albeit in conjunction with artificial landmarks.

How far are the algorithms of the computer vision community from enabling such a navigation system without environmental modifications- by recognizing natural visual landmarks? (Takeuchi & Hebert) demonstrates a vision system capable of discriminating several robot locales using image sequences. Such research in *place recognition* is sure to mature quickly, for the underlying algorithms have already been developed by the image retrieval community (IEEE 1998). We are hopeful that the visual tools that soon result will significantly lower the computational overhead of mobile robot navigation.

In conclusion, Sage is inspiring as an integration project, a quest for reliability, a human-computer interaction challenge and an educational endeavor. There is a popular stereotype that a project of this size forces the integration and reliability engineering efforts to overshadow all else. To the contrary, this project demonstrates that there is tremendous value in expending the energy to create long-term robotics applications.

Two important examples are born out by the Sage project. First, the existence of a situated, social robot led naturally to fruitful collaboration between engineers and experts in education and interaction. Roboticists do not have the training to solve social robotics problems alone; this paper aims to demonstrate a process by which roboticists can join forces with non-robotics experts to refine solutions at the formative level.

Second, the completion of a long-term robot application immediately opened new research avenues to the team; the long-term parameter adjustment research described above is an example of such research. Future plans for Sage also involve further use of the robot as a research platform.

Another research project soon to begin will imbue Sage with greater sensory and effectory richness to effect more meaningful interaction between the robot and museum visitors. The ultimate goal is to achieve two-way conversations between Sage and museum visitors. We look forward to the collaboration that will be required between The Robotics Institute and the museum's Divisions of Education and Exhibits in making this a reality.

Acknowledgements

Credit for the first tour guide robot goes to Ian Horswill for Polly the Robot (Horswill 1993). Jay Apt and Red Whittaker are the original authors of the idea to install a robot exhibit in the Carnegie Museum of Natural History's Dinosaur Hall. Nomadic Technologies, Inc. provided invaluable guidance on this experimental application of their XR4000. Parag Batavia, Salvatore Desiano, Marti Louw and Alvaro Soto all made useful comments on early versions of this article. Salvatore Desiano is responsible for the Sage schematic. Thanks also to David White of Mobot, Inc.

References

Borenstein, J. and Koren, Y. 1991. The Vector Field Histogram – Fast Obstacle-Avoidance for Mobile Robots. *IEEE Journal of Robotics and Automation*. 7(3): 278-288.

Burgard, W. Cremers, A., Fox, D., Hahnel, D., Lakemeyer, G., Schulz, D., Steiner, W. and Thrun, S. 1998. Experiences with an Interactive Museum Tour-Guide Robot. *Artificial Intelligence*, in print.

Burgard, W. Cremers, A., Fox, D., Hahnel, D., Lakemeyer, G., Schulz, D., Steiner, W. and Thrun, S. 1998. The Interactive Museum Tour-Guide Robot. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. Madison, Wisconsin. AAAI Press.

Canny, J. 1988. *The Complexity of Robot Motion Planning*. ACM Doctoral Dissertation Awards. MIT Press.

Kaelbling, L., Littman, M. and Cassandra, A. Planning and Acting in Partially Observable Stochastic Domains, *Artificial Intelligence*, Vol. 101, 1998.

Castellanos, J.A., Tardos, J.D., Schmidt, G. 1997. Building a global map of the environment of a mobile robot: The importance of correlations. *Proceedings of the 1997 IEEE Conference on Robotics and Automation*.

Drummond, M., Swanson, K., Bresina, J. and Levinson, R. 1993. Reaction-First Search. Proceedings of IJCAI '93. Chambery, France. Morgan-Kaufmann.

Drummond, M. 1989. Situated Control Rules. In *Proceedings of the Conference on Principles of Knowledge Representation & Reasoning*. Toronto, Canada. Morgan-Kaufmann.

Dugan, B. and Nourbakhsh, I. 1993. Vagabond: A demonstration of autonomous, robust outdoor navigation. In *Video Proceedings of the IEEE International Conference on Robotics and Automation*.

Elfes, A. 1987. Sonar-based real world mapping and navigation. *IEEE Journal of robotics and automation*. Vol. RA-3, No. 3. 249-265.

Gat, E. 1992. Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Control Real-World Mobile Robots. In *Proceedings of the Tenth National Conference on Artificial Intelligence*. MIT Press.

Hayes-Roth, B. and van Gent, R. 1997. Story making with improvisational puppets. In Proceedings of the International Conference on Autonomous Agents. Los Angeles.

Horswill, I. 1993. Polly: A Vision-Based Artificial Agent. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*. AAAI Press.

IEEE Workshop on Content-Based Access of Image and Video Database in conjunction with CVPR '98, June 21, 1998, Santa Barbara, California

Knotts, R., Nourbakhsh, I. and Morris, R. 1998. NaviGates: A Benchmark for Indoor Navigation. In *Proceedings of the Third International Conference on Robotics for Challenging Environments*. ASCE.

Kunz, C.; Willeke, T. and Nourbakhsh, I. 1999. Automatic Mapping of Dynamic Office Environments. *Autonomous Robots*. In print, Spring 1999.

Lazanas, A. and Latombe, J.-C. 1995. Motion planning with uncertainty: a landmark approach. *Artificial Intelligence* 76(1-2): 287-317.

Moore, A., Schneider, J., Boyan, J. and Soon Lee, M. 1998. Q2: Memory-based active learning for optimizing noisy continuous functions. In *Proceedings, The International Conference of Machine Learning*, Madison, Wisconsin.

Moravec H.P. and Elfes, A. 1985. High resolution maps from wide angle sonar. In *Proceedings, 1985 IEEE International Conference on Robotics and Automation*. Silver Spring, MD.

Murray, D., Jennings, C. 1997 Stereo vision based mapping and navigation for mobile robots. *Proceedings of the 1997 IEEE Conference on Robotics and Automation*. 1997.

Nilsson, N. 1996. Challenge Problems for Artificial Intelligence: Toward Flexible and Robust Robots. *Proceedings, Thirteenth National Conference on Artificial Intelligence*, pp. 1344 - 1345. AAAI Press. 1996.

Nourbakhsh, I., Powers, R. and Birchfield, S. 1995. Dervish, An Office-Navigating Robot. *AI Magazine* 16(2).

Roehrig, M. and Stockdale, M. 1998. *An Interactive Mobile Educator in the Carnegie Musuem of Natural History's Dinosaur Hall: Measuring the Effectiveness of Chips as an Educational Tool in the Museum Setting*. Program Evaluation, Carnegie Mellon University and Carnegie Institute of Art.

Schwartz, J. & Sharir, M. 1983. On the "piano movers" problem: II. General techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4:144-154.

Simmons, R. & Koenig, S. 1995. Probabilistic robot navigation in partially observable environments. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. Montreal, Canada. Morgan-Kaufmann.

Takeuchi, Y. and Hebert, M. 1998. Evaluation of Image-Based Landmark Recognition Techniques. CMU Technical Report CMU-RI-TR-98-20. Carnegie Mellon University. Pittsburgh, PA. 1998.

Thrun, S., Bennewitz, M., Burgard, W., Cremers, A.B., Dellaert, F., Fox, D., Haehnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. 1999. MINERVA: A Second-Generation Mobile Tour-Guide Robot. See www.cs.cmu.edu/~thrun/papers/index.html.

Thrun, S. 1995. An approach to learning mobile robot navigation. *Robotics and Autonomous Systems* 15(1995): 301-19.

Ushida, H., Hirayama, Y., Nakajima, H. 1998. Emotion Model for Life-like Agent and Its Evaluation. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. Madison, Wisconsin. AAAI Press.

Velasquez, J. 1998. When Robots Weep: Emotional Memories and Decision-Making. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. Madison, Wisconsin. AAAI Press.

Velasquez, J. 1997. Modeling Emotions and Other Motivations in Synthetic Agents. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*. AAAI Press.