

The Mobot Museum Robot Installations: A Five Year Experiment

Illah R. Nourbakhsh[§], Clay Kunz, Thomas Willeke

[§]The Robotics Institute, Carnegie Mellon University
Pittsburgh, PA 15213 illah@ri.cmu.edu

Abstract

This paper summarizes a long-term project to install and maintain socially interactive, autonomous mobile robots in public spaces. We have deployed four robots over the past five years, accumulating a total operational time exceeding seven years. This document introduces the robots, then focuses on lessons learned from each deployment. Finally, this paper describes how this entire project came to a close, offering a cautionary tale for those who wish to embark on such an effort in the future.

Introduction

The history of autonomous mobile robotics research has largely been a story of closely supervised, isolated experiments on platforms which do not last long beyond the end of the experiment. In January 1998, we and others began work on Chips, an autonomous robot intended to be more than a short experiment. Our goal would be to install Chips as a permanent member of the museum staff at the Carnegie Museum of Natural History in Pittsburgh, Pennsylvania [1].

Shortly thereafter, Mobot, Inc. was incorporated with a charter to improve and extend the Chips technology in a series of robot installations. Following Chips, three more robots were developed in succession; three of the four operated every day until February 2002. Of the four robots, three are museum robot installations, offering visitors of various exhibit spaces augmented detail regarding the exhibits at hand. Together, these three museum robots have logged more than 2,500 days of operation in separate, real-world public spaces.

In striving to deploy autonomous robots in a social niche, we had two high-level goals. First, the robots must be autonomous to the greatest extent possible. Human supervision of a full-time robot is unacceptable. At most, the robots should require only occasional human help, and should request such help explicitly. Even the routine trip to a battery charger should be performed by the robot autonomously.

Secondly, since the robots would be deployed in public, they must have sufficiently rich personalities to achieve compelling and fruitful interaction with humans in their

environments. Note that we care not just about 'compelling' but about 'fruitful' as well. These robots have an educational charter and are therefore justified only if they demonstrate real educational efficacy.

In the end of the day the robots did achieve some measure of educational efficacy as well as long-term robustness, but the social mission of the robots as well as the commercial justification for such robotic endeavours proved to be a challenge beyond our reach. These autonomous robots are now off-line, as of April 2002, and so, as the penultimate section of this paper explains, this long-term experiment is now complete.

Robot Overview

The three robots described in this paper all share the same basic motion platform (the holonomic Nomadic Technologies XR4000 base); the same operating system (RedHat Linux); and the same programming environment (Gnu C++). The first robot of the series, Chips, began work at the Carnegie Museum of Natural History on May 22, 1998. Chips operated exclusively in Dinosaur Hall, which contains the large bone collections of *T. Rex* and other massive dinosaurs as well as ancillary exhibits focusing on topics such as paleogeology and ancient aquatic life. Chip's charter was to provide tours in Dinosaur Hall, presenting audiovisual information regarding both the large bone collections as well as the less frequented, smaller exhibits. Until it was taken off-line, Chips operated for just under 4 years, covering a total travel distance exceeding 500 km.

The second robot, Sweetlips, conducted tours in the Hall of North American Wildlife, also at the Carnegie Museum of Natural History. This space is comprised of dioramas, where preserved wildlife specimens are shown in naturalistic settings. This portion of the museum has extremely low visitor traffic, so Sweetlip's charter was to attract additional visitors and to then bring the static dioramas to life using high-quality video footage of the wildlife in their natural habitat.



Figure 1: Chips the Dinosaur Hall Robot

Sweetlips operated beginning May 19 1999 and covered a total distance exceeding 185 km autonomously.



Figure 2: The Sweetlips robot near a diorama

The third robot, Joe Historybot, operated in the Atrium of the Heinz History Center. Its mission was to welcome visitors to this historical museum and to provide both information and a tour of various permanent exhibits



Figure 3: Joe Historybot in the Heinz History Center

placed throughout the atrium. Joe provided historical context in an entertaining multimedia format. The robot also provided tutorials on topics such as speaking English with a Pittsburgh accent. Joe began operations on July 8, 1999 and covered a total distance exceeding 162 km during its total period on-line.

Lessons Learned

The underlying goals of compelling interaction and maximal autonomy have remained constant throughout the creation of all three robots. However, each succeeding robot was the product of a complete re-design phase based on lessons learned from previous robots. Although some technical aspects remained unchanged, such as the programming language and robot mobility chassis, virtually all else evolved in an effort to improve both the autonomy and interactivity efficacy of each robot.

We are in the unique position of having an established trajectory of real-world interactive social robots. Studying the evolving lessons learned from these installations may prove useful in uncovering information that is valuable to future robot installations. In the following two sections we present such lessons learned, discriminating between the two top-level goals of providing maximal autonomy and producing maximally effective robot-human interaction.

On Robot Autonomy

The first requirement of a public robot is safety, both for the general public and for the robot itself. At the heart of the matter is the robot's method for avoiding collisions, which must be especially robust, since the robots operate



Figure 4: The fiducial docking landmark for Chips

without supervision. It is notable that the collision avoidance code on these robots is the least changed code over the course of their creation and installation. The robots use 48 ultrasonic range-finding sensors to detect both static and dynamic obstacles in the robot's path. This obstacle detection problem is simplified by both the density of sonar sensors on the XR4000 and by the structure of the navigable spaces specified to each robot. Navigation is performed within polygons that serve as "highways." The robot plans to achieve position subgoals that lie at the intersection of adjacent highways. Furthermore the highways are specified as conservative, open spaces within the museum's hallways and open

areas. Given this particular abstraction of navigable space, any short sonar reading can be treated as an unexpected obstacle, greatly simplifying the detection component of obstacle avoidance.

The obstacle avoidance algorithm is further simplified by several factors. First, the XR4000 robot is truly holonomic, constraining the avoidance algorithm to changes in 2D trajectory alone, fully decoupling these decisions from the robot's angle which is independently governed by the vision system to maintain visual sighting of the fiducials at the end of hallways. Second, the speed at which we chose to move these robots in such crowded environments was sufficiently slow that dynamics would not need to be considered for safe stopping and turning.

With these simplifications in mind, the obstacle avoidance system was a simple case-based chain, taking into account only the desired direction of travel and the instantaneous sonar data. The case-based obstacle avoidance technique achieves very high reliability due to its simplicity and diagnostic transparency [2].

Because of the limited accuracy of sonar at close range, the robots will occasionally become stuck when they approach a wall too closely. Given the infrequency of this failure mode (less than once per week), we feel the increased peace of mind due to conservative motion primitives was worth the price.

A second critical aspect of autonomy in our unsupervised application is the ability to detect failure and signal humans for help. Early in the development of the Chips platform, we began using pagers, which the robot was able to signal via electronic mail messages. Both active requests for help were implemented as well as regular, twice daily pages indicating healthy robot status, enabling us to verify liveness throughout the week. Basic failure detection is straightforward on Chips due to its sensors, which provide almost full observability into its own physical state in the environment. Encoders enable direct measurement of actual distances traversed; the vision system verifies continuously that the optical fiducials remain in the field of view; and analog inputs verify that the batteries' charge/discharge performance has not yet degraded.

Initially, Chips sent for help quickly, giving up as soon as a failure was detected. Soon we began adding diagnostic methods to reset subsystems that were not functioning correctly. This evolved into a generic *retry* method for diagnostics within our software architecture: each time a task is performed, check the result for validity. If the command failed, then reset/undo the device and try again. Surprisingly, this simple strategy has a commanding effect on a complex robot's failure rate. Of course, for each command that can be retried, an additional reset/undo method must be available. For physical processes such as docking, this can present kinematics challenges.

In order to achieve true self-reliance, each robot must be able to recharge itself when necessary. This is accomplished using a simple 3D fiducial, aligned with an electrical outlet that provides both translational and rotational position feedback (see Fig. 4). Using this marker, the robots have demonstrated reliable positioning to an accuracy of 1.5 mm using visual position servoing. The entire docking process, including moving over a distance of 4 m to the outlet and then fine-servoing for the insertion process, takes less than three minutes.

The *retry* method comes into play even in the case of this docking maneuver. If the battery voltage following docking fails to rise, the robot will physically reset by backing out of the plug and into the hallway. Then, it will repeat the docking attempt. This policy is effective in most cases because, although the code is apparently deterministic, there is sufficient nondeterminism in the environment that the overall system can have different outcomes when run consecutively.

This general approach is now used for entire classes of robot failure, including but not limited to: battery overcharging and undercharging exceptions; framgrabber anomalies; DVD player errors; encoder value errors; emergency-stop activation errors; etc.

Finally, a critical ingredient for autonomy across all of these mobile robots is the ability to navigate autonomously and with extreme reliability within the space that is served. Our goal in the implementation of these interactive robots was to achieve mean time between failure rates exceeding 2 weeks, and to do so with totally *unsupervised* operation. To this end, we chose to minimize the chance of navigation error by modifying the environment with custom fiducials. The fiducials provided both approximate range information (via height in the field of view) and precise angle information (via lateral position) during traversal of most navigable "highways." Chips made use of a set of high-contrast, high-saturation paper landmarks placed at the end of three of its four travel hallways. As robots were installed in various environments, the fiducial markers' complexity and expressiveness increased, including natural edge detection of window-wall boundaries and a variety of color and light fiducials. Furthermore, due to varying lighting conditions especially prevalent in the Heinz History Center (due to large, open windows), we added methods to enable tracking multiple fiducials simultaneously as well as *retry* techniques for re-acquisition of lost landmarks.

Such artificial markers are detectable with such high accuracy that each robot's position was represented as a single discrete position hypothesis. Not only was multi-hypothesis tracking obviated by this approach; but Kalman filter-based sensor fusion for position update was




Figure 5: Chips attracts visitors of all ages

also unnecessary because a single fiducial sighting is sufficient for complete localization.

In summary, Chips, Sweetlips and Joe Historybot were able to autonomously navigate public spaces for days at a time, charging themselves as they saw fit. This level of autonomy achieved Mean Time Between Failure values of between 72 and 216 hours, and only with great effort would MTBF ever climb beyond such values for any real duration. Remaining failures were eventually stochastic and unpredictable, a tire failing here, and a light bulb failing there. However, a significant achievement with respect to the project as a whole was that, following the first two years of effort, nearly every robot failure was autonomously detected by the robot itself. The days of robot failure unannounced by the robot itself were quickly over in the course of the experiment. The total autonomous travel distance for the combined set of robots exceeded 840 km.

On Interaction and Educational Efficacy

Our second requirement was to deploy robots with compelling interactivity. As the science of Human-Robot Interaction (HRI) is in its infancy, so it is not surprising that the robot interaction component was entirely redesigned with each successive deployment.

An interview with the exhibits maintenance staff of any large museum will drive home an important fact: people are basically destructive. Sometimes this purposeful damage is indeed caused by malicious visitors. More frequently, curious individuals who are trying to better understand the exhibits will cause damage. For example, some individuals attempt to push the robot off course physically to see if it will recover. Others will push any large red emergency stop button to see what happens next.

Also, what attracts people in a public space varies greatly depending on the context of a particular public space. When in an “entertainment” space, such as a museum, people will be curious and attracted by new and unusual exhibits. To that end the physical appearance of

a robot is extremely important as a visual hook. But two other characteristics produce even better results: motion and awareness. When the robot is in motion, it draws great attention from nearby people. To capitalize on this aspect of human behavior we found it useful to have some robots exhibit limited prosody during delivery of long, static presentations.

But the single most successful way for a robot to attract human interest is for the robot to demonstrate awareness of human presence. Interactions between humans and complex machines are typically initiated by humans. When a robot deliberately faces a person and says “Hello,” he or she is almost always both surprised and enthralled.

In contrast to entertainment venues, more utilitarian spaces such as large shopping centers and office buildings elicit far less pronounced reactions. In these spaces people tend to have an agenda; a schedule. They rush about from appointment to appointment and have little time to be side-tracked by new and entertaining machines.

One very important lesson learned from our experimentation is that attracting humans, while itself quite challenging at times, is far easier and very much unrelated to the skill of retention. Museum exhibit designers tend to make their exhibits more interactive, often even taking on the characteristics of a conversation, in an attempt to retain the visitor long enough to successfully pass on an educational nugget. An exhibit may pose a question requiring the visitor to lift a panel or push a button to hear the correct answer, for instance.

We have found that these techniques for audience retention are equally valid for HRI [3]. Chips presents long (two minute) video clips at various locations throughout its tour path. As our robots evolved, so did their level of interactivity. For instance, Sweetlips includes the human observer in the process of choosing a tour theme based on their interests. Joe goes further, answering many different classes of questions and even asking humans several questions in a game-like format.

Because of a robot’s particular sensory and effector strengths, dialogue is multimodal and not necessarily verbal. Thus, while the human may be pushing buttons or using a touch screen, the robot may be responding with spoken words, music, graphics, video, text, physical gestures and motion.

We have learned several lessons from such robotic dialogue design. Firstly, there often will be a crowd of people around the robot rather than a single person. Together with background noise from the environment, this will make it difficult or impossible for some to hear the robot’s responses if they are purely verbal. Therefore, all robot responses should be multimodal, including not only written text (e.g. captioning) but also graphics and video content.

Second, long presentations, even movies, are guaranteed to drive audiences away. Instead, short responses combined with interactive questions are most effective at extending the *time on task*. This parallels normal human-human interaction: the best conversations are dialogues, not monologues.

A final lesson learned with respect to HRI involves the psychological effect of creating anthropomorphic robots. There are strong social rules governing appropriate behavior between humans (though these rules vary between cultures and segments of society), and there are other behavior patterns that people follow when interacting with machines and computers. A robot that looks somewhat human and has a rudimentary personality falls somewhere between these modes.

The majority of people treat a robot as part human, part machine, clearly following some modified form of human interaction. Often they will treat the robot like a human by default, dodging its path and verbally responding to it naturally. If they become interested in some features of the robot, or want to investigate how it works, however, they will begin to treat the robot like a machine, ignoring social decorum by refusing to get out of its way and standing rudely in the robot’s path to elicit a reaction.

We theorize that humans use whichever social mode is most convenient for their short term goals. Fortunately, people will also often accommodate a robot that behaves in a socially unconventionally manner (were it a human).

A second avenue of exploration involves the use of affection in designing robot behavior. The main reason for a utilitarian robot to display emotion is that humans expect and respond to them in somewhat predictable ways. People have a strong anthropomorphic urge and tend to attribute internal state to anything that behaves appropriately. People are also conditioned to react to the emotions displayed by another person. These are powerful tendencies that robots could exploit.

These reactions are entirely behavioral. People cannot discern the true internal state of another human or robot. Their responses are thus entirely dependent upon perceived behavior. Chips and Sweetlips used sophisticated internal mood state machines that would change over the course of the day, affecting the behavior of the robot in the small and in the large. But since visitors to a museum only interact with a robot for a short period of time, the long-term mood shifts were moot.

For this reason, Joe Historybot uses no such internal mood representation and, instead, has a more transparent set of affective reactions to simple stimuli. For instance, stand in front of Joe unflinchingly and it would blurt out, “This isn’t the Parkway and you’re not PennDOT!”

The eventual goal of each interactive robot is to transfer information to human visitors. To test the educational efficacy of Chips, outside evaluators were invited to conduct an educational study [4]. The formal

focus of this project was to answer the question: *Is Chips 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.*

Table 1: Educational concept questions: success rates before and after robot tours

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

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 Chips 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 Chips for the same duration. However, 74% of visitors remained with Chips for between 5 and 15 minutes. The most significant differences between Chips and docents based on questionnaire results involved tour speed and sound level. Chips’ overall speed was viewed more favorably while docents were rated as much easier to hear. Interestingly, these results were consistent across both adult and youth age groups.

Educational efficacy was measured by asking adults and children knowledge-testing questions both before and after robot-led tours. The questions for adults and the success rates before and after robot tours are shown in Table 1. These results were pleasing in that they establish a quantitative educational efficacy for the robot tour guide.

In summary, the interactivity of our robots has evolved along four axes: engagement, retention, dialogue and anthropomorphic/affective qualities. Although this field of research is quite young, it is already clear that there remains great plasticity in the human-robot interaction model: human biases and bigotry regarding robots are not yet strong nor fixed. We have an opportunity, as robot designers, to guide, not just robot behavior, but the default

human behavior that will lead to the most fruitful possible human-robot interaction of the future.

The End of an Experiment

Although launched with much fanfare, this series of robotic experiments came to an end almost silently. Several factors came together almost simultaneously to bring about this end. Any museum is in the business of image. Add to this the extreme conservatism that is necessarily part of running most museums, and you have a formula for shunning unnecessary technological advances. That two robots were able to become docents at the Carnegie Museum of Natural History is itself somewhat miraculous and is due in no small part to the vision of Jay Apt, former astronaut and former director of that museum. With his retirement from CMNH, the robots, his brainchildren, were left without a champion among the museum executives.

Indeed, the educational study demonstrated real educational efficacy; this is a necessary but insufficient reason for a museum such as CMNH to spend more than \$10,000 each year paying for labor and parts (primarily new high-quality batteries every 3 months) to keep a 24-hour robot running.

The second deciding factor was purely economic. The museum market is small, with only a handful of museums that are large enough to purchase the hardware and audio-visual content for a \$200,000 robot system.

The interface trend that was proven during the progression from Chips to Joe Historybot is that greater bidirectional activity significantly increases time on task with an autonomous robot. In the end, this natural progression has led to a new educational robot that is not independent or autonomous, but rather a member of a human-robot tour-giving team. In the National Aviary (Pittsburgh, Pennsylvania), Mobot has produced the RAVEN robot, which provides a tour-giving docent with access to extensive multimedia and also serves as a mobile videoconferencing router so that school groups can have telepresent tours of the National Aviary. While the robot must still effect highly reliable obstacle avoidance, navigation is relegated to visually following the human docent while maintaining a safe distance from obstacles. Thus the autonomous robot that threatens to replace a docent becomes the docent's helpful aide.

Conclusions

Over the course of the past five years, we have built three museum tour guide robots that have each operated on a daily basis with the public, autonomously and without direct human supervision. While this has been done before [5,6,7,8] our robots are unique in their completely

unsupervised free-roaming obstacle avoidance, and in their mission to entertain and inform the generally public in documented, educationally effective ways. We have learned many useful lessons in attempting to meet the above challenges; perhaps the most striking is that it is indeed possible to deploy robots like these in public places, unsupervised, for long stretches of time. These autonomous robots are now off-line, after almost five years of operation, and so this experiment is now complete.

Acknowledgements

The original Natural History museum robot concept is due to Jay Apt and Red Whittaker. The following companies provided invaluable help: Nomadic Technologies, Magic Lantern, Maya Design, Redzone Robotics. Thanks also to co-creators: Judi Bobenage, Ben and Eva Carter, Vinton Coffman, Jan Falkowski, Sebastien Grange, Iacovos Koumas, Ronald Lutz, Roland Meyer, Carolyn O'Brien, Todd Simonds, Alvaro Soto and David White.

References

- [1] Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, R., Meyer, R. and Soto, A. An Affective Mobile Educator with a Full-Time Job. *Artificial Intelligence*. 114 (1-2), pp. 95-124. October 1999.
- [2] Nourbakhsh, I. Property Mapping: A simple technique for mobile robot programming. In *Proceedings of AAAI 2000*. July 2000.
- [3] Nielsen, J. *Usability Engineering*. Boston, MA: Academic Press, 1993.
- [4] Roehrig, M. and Stockdale, M. *An Interactive Mobile Educator in the Carnegie Museum 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.
- [5] Thrun, S., Beetz, M., Bennewitz, M., Burgard, W., Cremers, A.B., Dellaert, F., Fox, D., Haehnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva. *International Journal of Robotics Research*, 19(11):972-999, 2000.
- [6] Siegwart, R., Arras, K., Bouabdallah, S., Burnier, D., Froidevaux, G., Greppin, X., Jensen, B., Lorotte, A., Mayor, L., Meisser, M., Philippsen, R., Piguat, R., Ramel, G., Terrien, G. and Tomatis, N. Robox at expo.02: A Large Scale Installation of Personal Robots. Special issue on Socially Interactive Robots, *Robotics and Autonomous Systems*, 42 (2003) 203-222.
- [7] Sarcos Corporation. www.sarcos.com
- [8] Pyxis Corporation. www.pyxis.com.