Robots and Education in the classroom and in the museum:

On the study of robots, and robots for study

Illah R. Nourbakhsh
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
illah@ri.cmu.edu

Abstract

The Mobile Robot Programming Lab and the Toy Robots Initiative at Carnegie Mellon's Robotics Institute both focus on the interaction between humans and robots for the sake of education. This paper describes two specific endeavors that are representative of the work we do in robot education and human-robot interaction. The first section describes the Mobile Robot Programming Lab curriculum designed for undergraduate and graduate students with little or no mobile robotics experience. The second section describes the process by which an edutainment robot, *Insect Telepresence*, was designed, tested and evaluated by our group and CMU's Human-Computer Interaction Institute.

Introduction

Mobile robots ground ethereal computation to palpable physicality, and as a result they have the potential to excite and inspire children and adults in a way that no desktop computer can. Robots themselves make an inherently fascinating hybrid-dynamical system for study and comprehension. At the same time, they are also excellent facilitators for the study of fields as diverse as software engineering and the natural sciences.

In this paper, we describe specific projects that we have undertaken in the form of Robot for Study and Robot as Facilitator. At the educational level, we have designed and produced robotics curriculum for science centers, summer institutes for talented elementary students, and college-level students (Nourbakhsh, 2000a). All of this curricula shares a common approach, in which students work in teams in a problem-based laboratory environment. In this paper, we describe just one of these curricula, specifically the undergraduate/graduate *Mobile Robot Programming laboratory* course that has been taught for seven years at Carnegie Mellon and earlier at Stanford University.

The Toy Robots Initiative concentrates on creating robots that can be transparent facilitators, enabling humans to engage in experiences that are otherwise impossible. In the second half of this paper we describe the T.R.I.'s first installed artifact, *Insect Telepresence*. This discussion presents the techniques that we use to design and evaluate interactive robot-human systems.

Funded in part by the Vira Heinz Endowments, Pittsburgh, PA.

Robot for study: Mobile Robot Programming

Robotics Laboratory one decade ago meant, with few exceptions, building a physical robot, then writing limited low-level software for the system while continually debugging the mechanicals and electronics. However, with the advent of low-priced mobile robot chassis (e.g. Nomad Scout, IS-Robotics Magellan, ActivMedia Pioneer), a new space of possible robotics courses exists: mobile robot *programming* laboratories.

Mobile robot programming is a software design problem, and yet software engineering best practices are not very helpful. Mobile robots suffer from uncertainty in sensing, unreliability in action, real-time environmental interactions and almost non-deterministic world behavior.

Mobile robot programming faces sufficiently unique challenges that research projects have concentrated on creating languages and architectures that facilitate the task (Bonasso and Kortenkamp, 1996), (Brooks, 1986); (Firby, 1987), (Horswill, 1999), (Simmons, 1994).

Are mobile robots a truly unique target for learning not just software engineering but also problem-solving? The examples that follow are meant to convince the reader:

The physical world is virtually non-deterministic.
Although this may seem in jest, it is indeed the case that mobile robot programming can produce behavior that consists of long stretches of appropriate behavior, punctuated by rare moments of almost irreproducible error. Example: a student team creates a wandering program for a Nomad 150 robot and all is largely well. Sometimes, once every few hours, the robot smacks straight into a wall. All the rest of the time, performance is flawless. The reason turns out to be a rollover in the rotational encoder: when the robot is placed in its starting position just so, with bad luck to the tenth of a degree, then the error occurs. Otherwise, it is not reproducible. To the untrained eye, the robot and environment seem altogether non-deterministic.

Knowledge of physics is essential.
The environment is such a critical part of a mobile robot system that knowledge of the environment is underrated. Example: our students created maze-navigating robots one semester, and the final contest is held in a new conference room. The robots *all* fail consistently. The error: every five seconds, all the robots stop, as if synchronized. The problem: the robots use active infrared to recognize obstacles. The fluorescent lamps in the new conference room, unlike all the other fluorescent lamps in the building, produce infrared spikes every few seconds, saturating the robots' sensors.

The workspace can be changed too easily.
Mobile robots are so portable that, unlike most manipulator robots, factory robots and, indeed, desktop computers, their environment can be drastically changed with ease. The general result is that the unconscious assumptions of the scientist, which are usually left uncovered, are quickly brought out in the open when the mobile robot is taken from research lab to demonstration conference room. Example: a robot navigated by visually recognizing the floor. It also automatically adjusted its own saturation and hue thresholds. Unbeknownst to us, it had accommodated the very high saturation of the bright red carpet in our lab, ignoring all other cues. When taken to another environment, it failed altogether, because the only reason it was working properly in the lab was that the floor was far more saturated than every pair of shoes the researchers wore.

Errors cancel each-other out effortlessly.
Mobile robots are so complex that a hardware bug and a software error can actually compensate for one-another temporarily. This unusual phenomenon is often observed when the software is somewhat adaptive. Example: a tour guide robot, Chips, began to experience poor traction with one wheel because of a slipping tire (Nourbakhsh et al., 1999). The program increased the bias of its turns in one direction, compensating for this gradual degradation. When the tire finally broke off and was replaced, the robot was unable to track a straight line, and the reason was surprisingly hard to discover.

The perceptual stream is human-incompatible.
Mobile robots can have thousands of data points per second streaming in from a laser rangefinder, hundreds of data points per second from sonar and infrared, and tens per second from tactile sensors. The resulting instantaneous percept, not to mention the space of perceptual history, is far larger than that which a human can interpret effectively, particularly in real time. Example: Dervish, our office-navigating robot, has sonar on both sides to sense doorways and hallways (Nourbakhsh, 1998). During competition preparation, a sonar on one side of Dervish failed. However, its probabilistic navigation system continued performing using the single working side. Our only clue was its slightly degraded ability to collapse its uncertainty. But since the robot had been moved to a new building in a new city, the furthest thing in the roboticists' minds was that this slight degradation was due to the total failure of a hardware component.

Side effects show neither spatial nor temporal locality.
Mobile robots are exposed to varying environments, and certain environmental features may be seen extremely infrequently. As a result, an introduced bug in robot software may only affect mobile robot behavior once the robot is far away from the origination of the bug, both in space and in time. A telling example involves sonar firing frequency. On our Nomad Scout robots, the firing rate can be adjusted through software. When the students adjust the rate to a level that is extremely high, the robot looks as if performance has only improved, with values from the sonar boards that are far less stale. However, there are certain pathological configurations of MazeWorld that, at the higher firing rate, cause the sonar to interfere just so, resulting in a false short reading on the front sonar and, for most robots, false positive detection of an obstacle.

In summary, mobile robots provide a programming challenge that is unique: the environment, the robot hardware and the software itself all play equally important roles in the behavior of a mobile robot. The mobile robot programmer learns to understand real-time issues, experimental processes, some elements of radiation physics and of course the psychology of diagnosing a misbehaving robot. We developed the course *Introduction to Mobile Robot Programming* in order to expose students to this broad problem-solving exercise.

Introduction to Mobile Robot Programming

Introduction to Mobile Robot Programming is a one semester course that uses a problem-driven syllabus to transform novice roboticists into expert mobile robot programmers. Students form fixed teams of three and four at the beginning of the semester and immediately set out to address weekly lab assignments. The course outline, shown in Fig. 2, shows the progression of assignments.

Throughout the semester, students learn new techniques for transforming a sequence of sensory inputs to desired outputs. They begin at the lowest level and most real-time aspect of this transformation, implementing closed-loop controllers and functional/reactive algorithms for position control and obstacle avoidance. Then, they repeat the process but with abstract sensors and effectors (e.g. maze recognition sensors and navigation actions). The students not only end with highly effective navigation robots, but experience the power of hierarchy and abstraction, as well as the challenge of writing their own abstract percepts and actions.

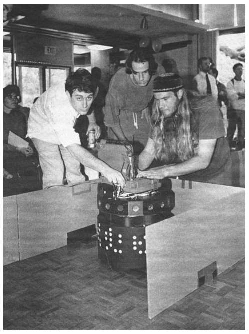

Figure 1: The 1999 Mobile Robot Programming Contest

Roughly halfway into the semester, the students transition from low-level motion and sensor interpretation code to AI-level planning code, designing a system that executes conditional plans, then designing a system that automatically generates and executes conditional plans.

0 Getting Started
0.1 Course description
0.2 The Robot Manual
0.3 A Symantec Java primer
1 Introduction: The Art of Robot Programming
1.1 Writing behaviorally transparent code
1.2 Playing sound in Java
1.3 Lab Handout: *The Dumb Wanderer, The Head Turner*
2 Feedback Control and Reactive Control
2.1 The Scout's motor controller unveiled
2.2 Functional reactive programming versus case-based reactive programming
2.3 Lab Handout: *The Position Commands, Run-Away, Smart Wanderer*
3 MAZEWORLD: Sensor Interpretation and Corridor-Following
3.1 Introduction to Mazeworld (new software distributed)
3.2 Operating characteristics of sonar sensors
3.3 Corridor-following and discretization
3.4 Lab Handout: *Sensor Interpreter, Corridor Follow*
4 GTNN: Creating a robust abstract action
4.1 Robot encoders and cumulative error
4.2 Localization
4.3 Lab Handout: *Go To Next Node*
5 Programmed Systems: executing plans and tracking state
5.1 States and state sets: formal concepts
5.2 Programmed systems: sequential and conditional plans
5.3 Universal plans
5.4 Environmental state-set tracking automata
5.5 Lab Handout: *Conditional Plans, Universal Plans, SSTA's*
6 Deliberate Systems: Planning
6.1 Search (DFID)
6.2 Sequential planning in state space
6.3 Sequential planning in state-set space
6.4 Conditional planning
6.5 Lab Handout: *Sequential planners, Conditional planners*
7 Architectures for Interleaving Planning and Execution
7.1 Programming architectures: balancing deliberation and execution
7.2 Assumptive and probabilistic planning
7.3 Heuristics for interleaving planning and execution
7.4 Lab Handout: *Assumptive Programming Architecture, Interleaved Prog. Arch.*
8 GAMEWORLD: The one robot game player
8.1 Introduction to Gameworld (new software distributed)
8.2 Lab Handout: *The one player Game*
9 Cooperation: The two robot game player
9.1 Introduction to mobile robot cooperation
9.2 Using the radio modems (new software distributed)
9.3 Cooperation strategies for The Game
9.4 Sample mazes for the Cooperation Game
9.5 Lab Handout: *The two player Cooperative Game*
10 The Final Contest
10.1 Introduction to the Final Contest
10.2 Basic game theory: cooperation versus competition
10.3 Sample games
10.4 Lab Handout: *The Final Contest*

Figure 2 **-** Curriculum outline for *MRP*

The final four weeks of the semester are spent on a rewarding application of all the skills gained throughout the semester: the teams create teams of mobile robots destined to compete with other robot teams in a giant maze world. This final robot contest, pictured in Fig. 1, is open to the university community and quickly becomes a popular annual event as well as a source of pride for the students involved.

Apparatus

Introduction ot Mobile Robot Programming (*MRP*) began as a course taught on Nomad 150 robots with Powerbook 170 laptops on board running Macintosh Common Lisp. The hardware has evolved to its current instantiation, Nomad Scout differential-drive robots headed by Toshiba Tecra computers running the Symantec Visual Cafe Java environment under Windows '95.

The students spend the first few weeks with the robots interacting in the unstructured world, then quickly move to a MazeWorld scenario, in which giant cardboard walls inflict a discrete structure upon the world that makes high-level navigation far simpler. This MazeWorld construction is critical in enabling the course to graduate to AI level programming assignments because it facilitates low-level navigation sufficiently for the students to achieve success.

Below we list all of the documentation that is provided for the course, in addition to the physical robots, laptops and cardboard walls. All of this material as well as Java source code is available from the course's web sites (Nourbakhsh, 2000b).

Lab exercises provide the exact specifications of the lab challenges and also provide much of the background material and hints that the students need.

Introductory handouts describe the robot hardware and software and provide step-by-step ways for the students to begin communicating with the robots.

Java introduction provides a primer to Java and to the Symantec programming environment used in the course. This primer is designed to take a student who knows any language well (e.g. C, Lisp) and bring them up to speed with Symantec Java.

Course outline provides a roadmap of the semester

Java programming environment is a software package written by us and used as the programming interface by all students.

Lecture outlines describe important topics to be covered in each class session, under the assumption that, once per week, a lecture class session meets.

Evaluation forms designed for each laboratory exercise provide scenarios and behavioral grading schemes for testing the robots.

Maze database offers a set of small and large MazeWorld mazes for the Final Contest and for practice runs.

In addition there are a number of *MRP* Java tools given to the students throughout the quarter, at appropriate times. Each contains Java source code and a document describing how to use the tool.

SampleBot is a sample robot interface project to get the student started

SampleSpeech is a text-to-speech synthesis project

SamplePlanner is an example depth-first iterative deepening planner

MazeEditor is a project for displaying, reading and saving MazeWorld mazes

SerialComm is a project demonstrating how to use the radio modems for robot-robot communication during the Cooperation Games and The Final Contest

GameEditor is a graphical tool for displaying and saving GameWorld mazes

MRP Conclusions

This course was inspired by CS222, Discrete Systems, an introductory AI course taught at Stanford University by Michael Genesereth. In turn, this course has developed curriculum that has been used to varying degrees by other institutions. Most recently, Martha Pollack has taught *MRP* at the University of Pittsburgh using identical hardware and a modified version of our curriculum.

The students who have graduated from *MRP* now number in the hundreds. More than four are now working on space robotics throughout the NASA system. A handful are working full-time in the research and industry robotics sectors. Each year, roughly 30% of the students request further mobile robot curriculum and thereby attend *Advanced Mobot Programming* in the following term. Of this group 50% go on to publish the results of their work at major conferences. Finally, in three years graduates of *MRP* have entered the AAAI National Robot Contest. In two cases, they achieved first place and in the third year they achieved second place. Significantly, in all of these cases, the students' first exposure to mobile robotics was *MRP* during the prior semester.

The mobile robot is a true situated automata, functioning in an environment that is as much a product of its own physicality, sensors and effectors, as the matrix surrounding it. This situated artifact not only has excellent potential as a focus of study, as in *MRP*, but it also has the power to provide humans with sensors and effectors that are otherwise inaccessible. In the next section, we describe the *Insect Telepresence* project, which is one project from the Toy Robots Initiative that aims to bring humans closer to the natural world by using robotics as a facilitator.

Robot Facilitator: Insect telepresence

The Carnegie Museum of Natural History (CMNH) has an invertebrate zoology department, featuring exhibitions of its collection of fascinating and exotic species of insects. A primary goal of the CMNH is to impart knowledge of the natural world to visitors. Hands-on and live exhibits in particular are used to encourage active exploration of the natural world, including non-human inhabitants of the planet. Currently, visitors only make a cursory examination of many exhibits because they are not provided with the appropriate tools to delve more deeply into the subject matter; they spend on average 5 seconds scanning exhibition contents. In an attempt to address the problem of shallow visitor investigation of exhibits, museums often provide docents or museum guides as a way to help visitors understand and gain meaningful insight into what is on display. However, it is not feasible for museums to provide a tour guide for every visitor and every exhibit.

Another interesting problem facing the CMNH, specifically the invertebrate zoology department, is that insects are fantastically complex organisms but most of this complexity is hidden from the naked eye. Most humans do not easily see the details of insect anatomy and behavior. This inability to see the insect structures is amplified by the fact that museum visitors cannot get close enough to the insect to examine it due to the physical constraints of display terrariums.

The Insect Telepresence project has developed a robotic tool that allows the visitor to enter the insect world and bridge the gap between inquisitive museum visitor and insect. We used human-computer interaction techniques to research how the robot might be used in the museum and suggest designs for the interfaces between the museum visitor, the robot and the insects. Following the fabrication of a physical prototype, we conducted two experiments to evaluate the control and design decisions and the success of the robot as a tool to enhance visitor exploration and investigation. The final robotic Insect Telepresence system is now a permanent exhibit at the main entrance of the Carnegie Museum of Natural History. As of September 1999, the Insect Telepresence robot has been interacting with Madagascan hissing roaches, humans and scorpions for more than five months.

Telepresence

Telepresence is often viewed as an advanced form of tele-operation, a technology that has been the subject of research for nearly fifty years (Goertz & Thompson 1954). The standard form of telepresence, telepresence in *space*, enables humans to interact with an environment that is spatially out of their reach. Often, this technology is proposed for applications involving environments that are hostile or unreachable for humans: outer space, deep recesses of the ocean, radioactive sites, etc (Weisbin & Lavery 1994; Draper 1995).

A more recent development in spatial telepresence is tele-embodiment, whereby a human achieves telepresence in a familiar or comfortable human environment. Remote surgery has long been an important target application of this technology because the surgeon may be physically unavailable at the locale where surgery is required (Green et al. 1995). Embodiment for the sake of visual and auditory presence in a remote environment has been demonstrated by Paulos & Canny in the creation of Personal Roving Presences (Paulos & Canny 1998a, 1998b). Recently, our group has added this level of personal embodiment to a full-time robotic tour guide at CMNH (Willis et. al, 2000), enabling web visitors to tour the museum without actually travelling to Pittsburgh, Pennsylvania.

This project has a different focus: telepresence in *scale*. Our intention was to use robotics and optics technologies to offer visitors a new sense of scale, enabling exploration of an environment from the perspective of a small insect. The robot, therefore, acts as the embodiment of the museum goer, who gets to interact with insects on their own terms.

The Insect Telepresence project offered both HCI design challenges and robot hardware and software design questions: how should users manipulate the robot, how should the captured image be displayed, how *transparent* should the robot should be in the exhibition (should it be seen or simply enhance the user experience without making its presence known?), and should we protect the insects from the robot, and by extension, the museum visitor?

In this part of the paper we will outline how the robot was built, what HCI experiments were conducted to answer these questions, and how those results have informed a new robotic exhibit that is now in daily use at the Carnegie Museum of Natural History.

Initial HCI Research

This project relies heavily on technology, yet its underlying goal is human: to help people have a more meaningful museum experience. The field of Human-Computer Interaction provides formal techniques for addressing the interfaces between humans and technology. Current HCI practice focuses on user-centered design, which suggests that understanding the user and the tasks to be accomplished can lead to better design of the tool. A team of six HCI Masters Students at Carnegie Mellon University employed three HCI techniques to generate recommendations for Insect Telepresence and evaluate the ensuing prototype. This section summarizes the research and findings of the HCI team.

Contextual Inquiry

Contextual Inquiry and Design is a technique developed by Hugh Beyer and Karen Holtzblatt to understand how people use tools to accomplish tasks, and from this understanding to develop designs for building better tools (Holtzblatt & Beyer 1996; Beyer & Holtzblatt 1998). The specific aims of Contextual Inquiry and Design are to understand the task, the task environment and the user of the tool being designed. In this case we needed to understand how museum visitors use the invertebrate zoology exhibits, and how they use museum tour guides to enhance their experiences.

We conducted Contextual Inquiries with three staff members at the CMNH to get a broad understanding of the typical museum visitor and how they interact with the exhibits. We spoke with a teen docent in the invertebrate zoology division; a tour guide leading a *bug tour* for a birthday party; and the head entomologist for the Division of Invertebrate Zoology.

The teen docent inquiry was approached with three target questions:

1 *What kinds of interactions occur between visitor and insect?*
2 *What kinds of interactions occur between visitor and docent?*
3 *What kinds of interactions occur between docent and insect?*

Our belief was that understanding these interactions would aid in the design of a robot and an environment to facilitate visitor experiences in the same way that docents do. The interview and observation of the teen docent led to five conclusions:

1. *People have short attention spans, especially when reading label copy in an exhibit. People browsing spend approximately 5 seconds looking before moving on.*
2. *Opportunities to personally contact and examine the insect significantly lengthens the total interaction time.*
3. *Docents provide a more interactive and engaging experience, causing visitors to spend more time looking and learning.*
4. *People prefer live exhibits to displays of pinned insects.*
5. *The separation of visitor and insect by exhibition casing and glass makes it difficult for visitors to see bugs closely.*

In hindsight these results are intuitive. An important secondary conclusion was that the effectiveness of the Insect Telepresence robot may be quantitatively measured by timing the average interaction time between visitors and the exhibit.

The second contextual inquiry was a *bug birthday party*, a special event offered by the museum, where a child celebrating a birthday can bring friends, take a tour of the museum focusing on insects and, for the finale, learn how to pin a grasshopper professionally. The purpose of conducting this observation was to note how docents deal with groups of visitors, and how groups of visitors interact with exhibits. This observation led to a confirmation of one of the earlier conclusions and the addition of two new conclusions:

1. *Docents provide a more interactive and engaging experience, causing visitors to spend more time looking and learning. (a confirmation of an earlier observation)*
2. *People are simultaneously drawn to and repulsed by the insects.*
3. *Kids were hesitant to touch initially, but it became a cool thing to do and they spent a great deal of time simply looking at the insects up close when the glass barrier was removed.*

The final contextual inquiry consisted of an interview with John Rawlins, the head of the Division of Invertebrate Zoology. The interview was conducted to understand insect psychology issues, museum staff issues, and past experiences with technology and exhibitions. The department was aware of the potential for using technologies such as robotics in their exhibits and, importantly, these individuals had strong feelings about the potential negative impact of such technology.

As a robotic community considering installation of robot technologies throughout society, we must be keenly aware of these anti-technology feelings and their very valid foundations. After all, the final goal of human-robot interaction in this case is to improve the experience of the staff and visitors at CMNH.

Three conclusions were drawn from this inquiry:

1. *Technology for technology's sake is not useful. It gets in the way of the point of exhibitions, which is to learn about the natural world.*
2. *Technical exhibits lose their appeal for visitors quickly if the exhibit lacks a real experience.*
3. *People come to a natural history museum to see the natural world, not technology.*

Our conclusions, on the basis of the three contextual inquiries, were:

1. *There is a need to increase the time spent looking in an exhibition. Visitors do not spend very long looking at exhibitions where much effort to interpret on the part of the user is necessary*
2. *People like assistance when looking at an exhibit. Docents facilitate the museum experience and cause people to spend more time looking.*
3. *People prefer live exhibits to the pinned ones. Opportunities to examine things personally and more closely increases time spent with an exhibit.*
4. *People come to the museum to learn and have meaningful experiences, and need a robotic tool to enhance that experience, not interfere with it.*

Modeling

After conducting Contextual Inquiries, the team of HCI students proceeded to the next step in the Contextual Inquiry and Design process: modeling the environment. The team conducted an interpretation session to revisit the interviews and record what was learned in the form of pictorial models. Five types of models were created. *Workflow* describes the communication necessary to complete work. *Sequence* charts the actual steps necessary to complete a task. *Artifact* depicts the tools necessary to complete a task. *Culture* represents the unspoken aspects of completing a task, including policies, attitudes and values. *Physical* shows the physical layout of the work space

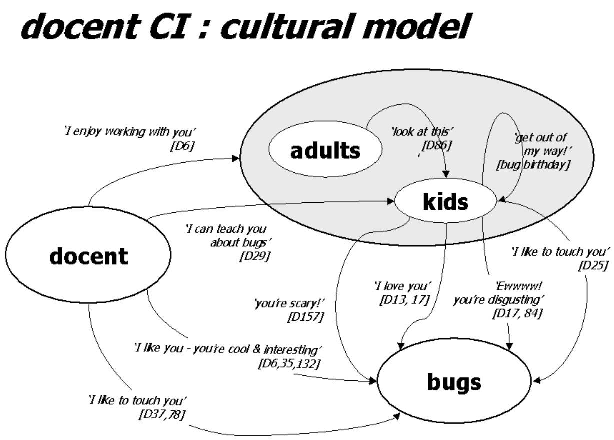

Figure 3: The teen docent Cultural Model

The five models give five different perspectives of the task, and make it possible to comprehend the complexity of the task. The models are used to develop a common language between the various stakeholders of the project. Graphical representations of this common language reveal the patterns and structure of work in a far more concise fashion than would be possible through prose.

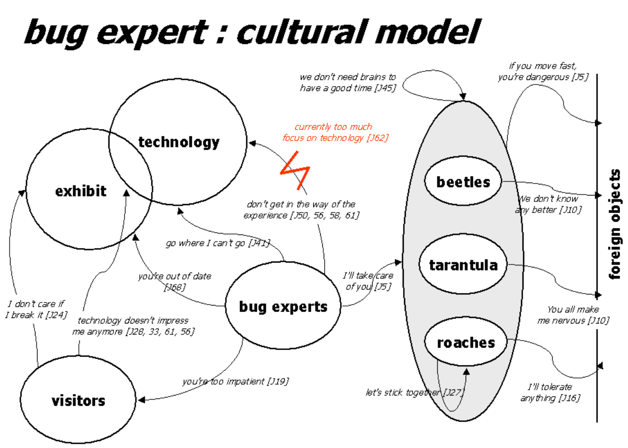

Figure 4: Expert cultural model

Figures 3 and 4 depict the two *Cultural* models based on the teen docent study and expert interview. For instance, in Figure 3, the docent's attitudes toward the insects are wholly positive while the visitors' attitudes include elements of both extremes. Much of the success of the docent tour can be attributed to the docent's ability to modulate these visitor attitudes, reinforcing visitors' positive attitudes while filtering out negative attitudes. This is just what the robotic exhibit must also accomplish.

Figure 4 depicts both the attitude of various insects at the possible existence of a free-roaming robot in their midst and the attitude of the bug expert with respect to technology. It is important to note that, although the bug expert is aware of the potential negative consequences of technology, he is also well aware of its advantages (e.g. *go where I can't go*). The modeling process was invaluable in transforming the bug expert from a hesitant member of the team to a champion for Insect Telepresence throughout CMNH.

Heuristic Evaluation

The HCI team employed a third technique for evaluating the usability of human-technology interactions to develop design solutions: Heuristic Evaluation. The technique requires usability or domain experts to evaluate an interface to judge its compliance with recognized usability principles, or heuristics. The principles used in these evaluations were developed by Jakob Nielsen (Nielsen 1994; Nielsen 1993). Two examples of recognized usability heuristics are:

> *Visibility of System Status*: The system should always keep the users informed about what is going on, through appropriate feedback within reasonable time.
>
> *Match between system and the real world*: The system should speak the user's language, with words, phrases, and concepts familiar to the user, rather than system-oriented terms. Follow real-world conventions, making information appear in a natural and logical order.

Each usability expert reviewed the interfaces independently, and the results were gathered and merged into one report. The report also reflected the team's observations with regard to the frequency and criticality of the heuristic violations found. The usability study was conducted by evaluating a technology in place at CMNH that most closely resembled *Insect Telepresence*. The conclusions of the study could then be used when designing and fabricating the robot exhibit.

The chosen technological artifact was a surveillance camera remote control. The camera, installed on the third floor of the museum, can be controlled from a kiosk on the first floor. The third floor camera image is projected on a large screen television, with four buttons on the face of a stand to control the rotation of the camera and the focus of the camera. In museums, there is an exponential decrease in visitor attendance with each successive, higher floor of the building. The purpose of this kiosk was to attract visitors to the third floor artifacts so that they would take the time to travel two floors up.

The conclusions of the heuristic evaluation were:

1. *The kiosk buttons did not readily communicate their use, making the device unintuitive and difficult to use.*
2. *Because the interface was not intuitive, they needed to have labeling or directions for their use. Few or no directions were provided to assist the user in understanding how to interact with the device, or what steps were needed to be successful in using the device.*
3. *Four buttons in a row to control left and right rotation and zoom in and zoom out functions do not match the user's idea of how to operate a camera.*

The lessons learned for application to *Insect Telepresence* are obvious: signage and input devices must be chosen with care in order for visitors to feel readily comfortable with the robot system. Since users do not have significant preconceptions concerning control systems necessary for interacting with robots, particular care must be taken in the first design of such an input device.

The Robot

Based on the HCI formative evaluations, it was clear that good human interface design for the Insect Telepresence robot would be critical to its success. But a second, equally important observation was that users would try to physically manipulate the environment in order to hit the insects. The final goal was to design a robot that would require no supervision during operation. In summary, the key design requirements were long life, an easy to use input device, robustness to user abuse and robustness to user attempts to damage the insects and robot.

At the same time, a number of optical solutions were considered to bring high-fidelity images of the terrarium to the user station. Laparascopic equipment, although of high fidelity, requires extremely high-intensity lighting and has a prohibitive price. Popular single-board cameras suffer both from poor lumens sensitivity and coarse pixel resolution. The resolution of the CCD chip would be critical because of the effective magnification caused by projecting the image onto a 35" video monitor.

Figure 5: The Toshiba micro-camera remote head and the robot raster

The Toshiba remote head camera system was chosen for its extremely small head size (5 mm x 30 mm) and lumens sensitivity (Fig. 5). The camera head houses only the lens and CCD chip. The digital signal is carried through a tether to a NTSC conversion box, which produces a video signal appropriate for the monitor and provides auto-gain and auto white balance controls.

Once the camera was chosen, a tether was clearly needed, obviating the challenge of designing a standalone robot capable of walking or rolling on the terrarium surface. Furthermore, placing the camera above the insects, canted down, would provide the illusion of navigating at their level while guaranteeing that no insects would be squashed. The robot that resulted from these design considerations is a XY-type raster with a camera mast that hangs down from above (Fig. 5). Camera angle in the *theta* direction would be controlled by rotating the entire camera mast from above.

In order to afford both high reliability and ease of service, drive motors were chosen as quarter-scale, internally geared high torque hobby servos (Hitec, Inc.) normally used in large-scale remote-control sailboats. These servos are designed for approximately 120 degrees of total travel. To overcome this obstacle, the X and Y servos were disassembled to remove the rotation stops and disconnect the feedback potentiometer. This transformed the X and Y servos from standard position-control devices to velocity control devices. Any offset from the current potentiometer setting would cause the servos to rotate continuously, with a small amount of speed variation depending on the disparity between the commanded potentiometer position and the current potentiometer position.

Because the camera mast housed the video tether without the use of a slipring, it was critical that the rotation servo be limited to a finite number of turns in each direction. This was accomplished by removing the internal feedback potentiometer in the rotation servo, then replacing it with an externally mounted, 3-turn potentiometer. Thus, the rotation servo remained a position control servo, albeit with far greater range of motion (1080 degrees) than originally intended.

Although the robot mast is intended to be manually positioned high enough to just clear the insects, a failsafe mechanism to avoid damaging the camera and insects would still be required, since the insects will at times climb the glass walls of the terrarium and climb on one-another. The solution was simple, based on the fact that the camera mast is a fairly long lever arm, exerting a significant tilting force on the raster assembly. This tilting force would disable low-friction wheels from engaging the tracks, and so the wheels were machined smooth. The smooth lucite surface of the wheel engages the steel tracks with just enough force to carry the assembly; any force exerted on the camera stops the entire robot in place.

The final requirement for robot design and fabrication involved control. The servos were commanded by an SV203 interface board (Pontech, Inc.) commanded by a standard 68HC11 board (Midwest Micro-Tek, Inc.). The microprocessor reads the joystick input values and the camera's angle, then computes a coordinate derive speeds to command the X and Y servos for camera-relative motion. In other words, *forward* on the joystick always corresponds to forward motion of the camera head by sending appropriate commands to the X and Y servos of the raster frame.

The basic input-out control and coordinate transform loop was written in C, compiled and placed on an EPROM on board the 68HC11, enabling museum staff to reset the software simply by power cycling the system.

Museum Experiment

A steel version of the Insect Telepresence robot was fabricated for installation at the Carnegie Museum of Natural History's main entrance area (see Section 4 for robot details). Museum grade exhibition cabinetry was designed, constructed and installed to house the robot and the insects, and to provide an environment for the Insect Telepresence experience, following a conclusion from the Think Aloud Study.

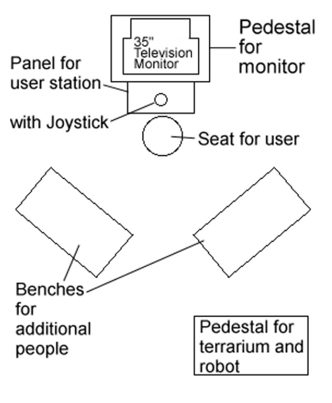

Figure 6: Overhead drawing of the display and cabinetry setup

The user station was constructed to place the visitor directly in front of (24-34" distance) a 35" high-quality color monitor which receives the camera image. A three-axis electric wheelchair joystick was installed, providing X and Y deflection on a single stick, topped off by a spring-loaded rotating knob at the tip of the joystick. The rotation control was used as a velocity controller of the camera's rotational position, as this knob had return springs and could not be used for position control.

People often come to the museum in groups, so benches were provided to accommodate and encourage spectators for the installation. The robot and terrarium were placed behind the viewer for two reasons: (1) to encourage concentration on the screen without the distraction of the sound and movement of the robot, and (2) as a way to visually draw people into the exhibit. The robot is visually engaging; many curious visitors were taken by its mechanical complexity, then moved on to the driving station to drive the robot.

Figure 7: Photograph of the Insect Telepresence installation

After installing the new robot and supporting hardware, user observations were conducted on three afternoons, over a period of seven hours. Users are defined as anyone who interacted with the installation, whether manipulating the robot, or simply watching the video display or vitrine. If a person participated in the installation for a period of at least 10 seconds, he or she was characterized as a user.

At the time of the observations, there was no signage or directions for use of the installation. The goal of monitoring museum visitors as they use Insect Telepresence was to discover if people could understand how to use the robot without assistance, and to take note of what people's natural inclinations are for use of the robot. Measurements and observations follow.

<u>Number of Users</u>

The total number of users observed: 204
Number of single users: 41 (~20%)
Number of users in groups: 164 (~80%)
Total number of user groups: 51
Average group size: 3 users

Most of the people engaging Insect Telepresence did so in groups. This can be attributed to how people visit museums, or that the age group that uses the exhibit tends to visit museums in groups. An interesting facet of the Insect Telepresence robot is the fact that it is visually appealing both at the control station and at the vitrine, where the robot and insects can be seen from above. This naturally leads to group activities, where members of the group work together to use the controls and observe the results at the vitrine.

User age

Average age of users: 19.5 years
Three modes: 8 years, 10 years, and 35 years

Although the average age of the user indicates a late teen, a histogram would indicate three modes in the range of users, with the ages skewed to the younger segment of users. Much of this can be attributed to the particularly high attention given to the exhibit by children who are accompanied by parents. A fascinating observation was that the Insect Telepresence robot engaged both the very young children, who immediately controlled the robot motion from the joystick console, and the young parents, who would also take a turn at piloting the robot.

Time on Task

Average time on task of all users: 60 seconds
Average time on task of a single user: 27 seconds
Average time on task for user groups: 93 seconds

These results are pleasing in comparison to the five second average time that visitors will often spend inspecting an exhibit. The Insect Telepresence exhibit engages visitors successfully. Note that the single user spends decidedly less total time in the installation than a group of users. Taking turns with friends can account for the difference in use time, as can play and exploration patterns based on working with other members of the group.

Additional Common Tasks

Looking between the monitor and vitrine while moving joystick: ~55 users (~27%)
Looking at people on the Monitor with camera: ~39 users (~19%)
Constant motion while driving: ~42 users (~21%)

There were a few common tasks that prevailed across all the users. All users were able to make use of the exhibit's primary mission, navigating using the image on the monitor as their visual guide to explore the insects. However, about 27% of those controlling the robot also looked at the vitrine during robot motion. Our hypothesis is that a person likes seeing that he is affecting the physical robot system with his control input. The vitrine is far enough away and the terrarium is high enough that users cannot navigate by looking back (indeed, the joystick's reference frame is camera-based, and so physical navigation by viewing the robot is extremely difficult).

Another common action was to keep the robot in constant motion while in control. About 21% of those using the Insect Telepresence robot view kept the robot in constant motion. Our theory for this action is that the users are engaged by the motion. The color monitor is extremely large, and so continuous motion is visually stunning to the point that motion sickness is possible. Those behaving this way seem to spend little time actually looking at the bugs, and more time driving like a video game.

Using the robot to look outside the terrarium at friends was another popular activity. Several groups shared the tasks of aiming the camera to look at friends and making faces for the screen.

Of those people who were alone when using Insect Telepresence, 41% participated by looking at the vitrine and watching other people use the robot. The remaining 59% of the single users also engaged in driving the camera.

Other Observations

Many users said aloud that the joystick controls the camera in the vitrine. There were ten instances where users verbally made the connection between moving the joystick, seeing the image change on the screen, and the movement of the robot in the vitrine.

Visitors were clearly engaged by Insect Telepresence, spending a significant amount of time studying the insects and interacting with the robot. Of particular interest is the connection made by visitors between their actions at the control station and the motion of the robot system in the vitrine. Recall that there was no signage hinting at this control aspect; and, furthermore, the control station and the vitrine were separated by approximately fifteen feet. This form of discovery was especially pleasing, as the visitors would grasp their potential for exploring the insect world and, at the same time, would appreciate the manner in which a technological artifact, the robot, was aiding them in this journey.

The Insect Telepresence robot has been operating continuously as of May 1999. It has operated in excess of 1000 hours as of September, 1999. The robot hardware has suffered one failure: the second rotation gear slipped vertically, requiring adjustment and tightening. The joystick, however, has been destroyed three times by visitors who deflect it with a great deal of force. After the third such event, the single, 3-axis joystick was replaced by two arcade-quality 2-axis joysticks that are designed to withstand roughhousing. Since that replacement was made, there have been no further system failures.

Conclusions

The Insect Telepresence robot is a successful mechanism to help museum visitors engage with the living exhibits at the Carnegie Museum of Natural History. Telepresence allows humans to enter the small-scale world of insects. The robot is a useful tool for learning more about insects by removing barriers to actually seeing what the insects look like and increasing the time spent looking at the exhibit. In addition, observing the insects in their daily life creates empathy for other life forms.

The robot allows the visitor to act as his own tour guide by offering the visitor control over what he sees. As this exhibition continues, and suggested educational materials are added, the visitor will be even better able to understand what he is looking at, and tailor the investigative experience to satisfy his curiosity about the fascinating and complex world just out of sight. The better able the Insect Telepresence robot is in supporting the visitor experience of examining and learning about the natural world, the more successful this venture will be.

References

Beyer, H. & Holtzblatt, K. 1998. *Contextual Design: Defining Customer-Centered Systems*. Morgan and Kaufmann, San Francisco, CA

Bonasso, R. and Kortenkamp, D. 1996. Using a layered control architecture to alleviate planning with incomplete information. In *Planning with Incomplete Information for Robot Problems*, AAAI Spring Symposium Series.

Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14-23.

Cox, M. 1994. Robotic telescopes: An interactive exhibit on the world-wide web. In *Second International World Wide Web Conference.*

Draper, J. 1995. Teleoperators for advanced manufacturing: applications and human factors challenges. In *International Journal of Human Factors in Manufacturing*.

Firby, R. 1987. An investigation into reactive planning in complex domains. In *Proc. AAAI-87.*

Goertz, R. & Thompson, R. 1954. Electronically controlled manipulator. *Nucleonics*.

Green, P., Hill, J., Jensen, J. & Shah, A. 1995. Telepresence Surgery. In *IEEE Engineering in Medicine and Biology Magazine*.

Holtzblatt, K. & Beyer, H. 1996. Contextual Design: Using Customer Work Models to Drive Systems Design. *Proceedings of CHI '96.* New York, NY: ACM.

Horswill, I. 1999. Functional programming of behavior-based systems. In *Proc. IEEE International Symposium on Computational Intelligence in Robotics and Automation*.

Nielsen, J. 1994. Heuristic Evaluation. Chapter 2 of *Usability Inspection Methods* Edited by Jakob Nielsen and Robert L. Mack. John Wiley and Sons, Inc., New York.

Nielsen, J. 1993. *Usability engineering*. Boston, MA. Academic Press.

Nielsen, J. 1992. Evaluating the thinking aloud technique for use by computer scientists. In H. R. Hartson & D. Hix (Eds.), *Advances in Human-Computer Interaction*, Volume 3, 75-88. Norwood, NJ: Ablex Publishing.

Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, R., Meyer, R. and Soto, A. 1999. An affective mobile robot educator with a full-time job. *Artificial Intelligence* 114:95-124.

Nourbakhsh, I. 2000a. Pointer to course material for elementary student robotics (C-MITES) and college-level robotics (undergraduate and grad). www.cs.cmu.edu/~illah/#Courses.

Nourbakhsh, I. 2000b. Pointer to source documents and source code for *MRP*. www.cs.cmu.edu/~illah/rix62.html contains all exercises and www.cs.cmu.edu/~illah/ROBOCODE contains all Java and C interface code and drivers.

Nourbakhsh, I. 1998. Dervish. *Artificial Intelligence and Mobile Robots*, Chapter 3. David Kortenkamp, R. Peter Bonasso, and Robin Murphy, eds. AAAI Press/ The MIT Press, 1998.

Paulos, E. & Canny, J. 1998a. Designing Personal Tele-embodiment. In *Proceedings of the IEEE International Conference on Robotic and Automation 1998*.

Paulos, E. & Canny, J. 1998b. PRoP: Personal Roving Presence. In *Proceedings of SIGCHI 1998.* ACM.

Simmons, R. 1994. Structured control for autonomous robots. *IEEE Transactions on Robotics and Automation*, 10:1.

Weisbin, C. & Lavery, D. 1994. NASA rover and telerobotics technology program. In *IEEE Conference on Robotics and Automation Magazine*.

Willis, H., Willeke, T., Apt, J. and Nourbakhsh, I. 2000. *Jurassic Classroom.* Pointer to the website for tele-embodiment of the Chips tour robot: www.clpgh.org/cmnh/jurassic.