Wiimote Interfaces for Lifelong Robot Learning

Micah Lapping - Carr
Odest Chadwicke Jenkins
Daniel H. Grollman
Jonas N. Schwertfeger
Theodora R. Hinkle
Department of Computer Science
Brown University
Providence, RI, USA 02912-1910
`mlapping|cjenkins|dang|js|thinkle@cs.brown.edu`

Abstract

We believe that one of the major impediments to involvement in the field of robotics and AI is the difficulty end-users face in creating viable robot control policies. We seek to address this difficulty with lifelong robot learning and have developed an intuitive robot control interface using the Nintendo Wii remote to aid in this task. From three large public demos and several smaller ones, we have gotten a multitude of positive responses on the interface. We also believe that others can find similar successes in the field of HRI using undergraduate researchers.

Introduction

I was nearing the end of my sophomore year and eagerly trying to find a summer internship. I had limited experience with robotics, having worked with the Robocup team for a semester, yet found the topic fascinating. Therefore, I asked the professor in charge if he would hire me for the summer. He did, and I began on the road that has brought me to where I am today, authoring my first paper. My name is Micah Lapping - Carr, and this is my story of how I joined and embraced the world of AI and robotics.

The state of robotics is constantly changing, but there is one barrier that will continue to impede its success if not addressed. For most commercial robots, only the "technically elite" (programmers and engineers) are currently able to create the robot control policies they want, while the rest of the population must make do using the built-in policies (such as those on the iRobot Roomba, or WowWee Robotics' line of Robosapiens) included by the robot's creators. Through **lifelong robot learning**, we aim to provide users of consumer robot technologies with a medium for transforming desired behaviors into robot control policies. Specifically, given the same situational awareness, a robot should make a decision similar to the one the creator of the policy would make. While several paradigms exist for such policy transfer (e.g. continuous teleoperation, speech and gesture-based instruction, text-based and visual computer programming, optimization/search), we remain confronted by a **human-robot divide**. This divide refers to the disparity between the needs and ideas of users in society, a population with a diverse set of technical abilities and creative design sensibilities, and their ability to instantiate robot control to meet their desired ends. If a **personal robotics revolution** is to come, similar to that of personal computing, there will need to exist applications that will make new forms of personal expression tangible, enhance personal productivity, and put this new technology into the hands of users (analogous to the spreadsheet, web authoring, 3D virtual worlds, etc., on the PC).

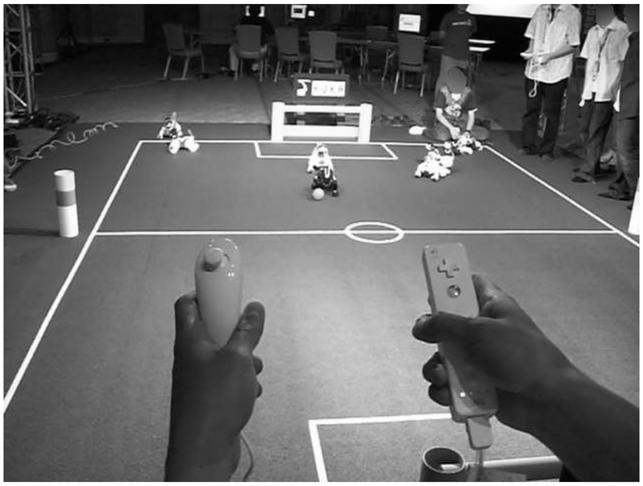

Figure 1: Nintendo Wii Remote and Nunchuk controlling a Sony Aibo playing robot soccer.

We believe it is crucial to provide as many ways as possible for people to experience and contribute to robotics. Furthering the analogy to personal computing, the computational sciences expanded over time as researchers and systems engineers solidified the foundations for modern computing. This expansion created many new pathways for application developers, such as high-level languages and a plethora of libraries. These foundation pathways continue to exist and are built upon by newer generations, but they are now complemented by pathways for design-oriented content creators (such as web development, graphic design, and an-

---

imation software), and end consumers (through office software, email, video games, etc.). In contrast, modern-day robotics focuses on the research and systems-engineering pathways due to the highly challenging nature of perception and action in the real world. With the rise and continued momentum of commercial robotics (typified by rising companies such as iRobot, WowWee, and Evolution Robotics, who are now being joined by more traditional companies such as Microsoft and Hasbro), new pathways for larger populations are beginning to appear, but they are not yet fully formed.

In the future, we envision *off-the-shelf* robots that can be "taught" by humans to perform **unknown tasks**, where no task-specific information needs to be hardcoded into the robot's decision making procedures. For instance, a user should be able to purchase a robot platform (let's say a robot dog) from their local electronics store and, without writing a single line of code, teach it to play soccer at a level competitive with hardcoded decision making. This makes the reasonable assumption that non-technical users are comparable or better creators of robot controllers, once technical barriers are relaxed. As examples, consider animated filmmaking, web design, and desktop publishing: while technical researchers and systems architects enabled the development of these media, it is often the design-oriented (artists, film-makers, etc.) that can make the most of these outlets for expression, in terms of aesthetics and accessibility. We believe the personal robotics industry will follow a similar route, in that once robot learning has been fully enabled by researchers such as ourselves, it will be the end users of robots who will be creating the best robot controllers, not the researchers.

Towards "human-guided" pathways into robotics, our aim is to realize lifelong robot learning through developing **scalable policy-learning algorithms suited for long-term human-robot interaction (HRI)**. While differing in methods, our objectives for lifelong robot learning fall under the broader scope of "socially-guided robotics" proposed by Breazeal and colleagues (Thomaz & Breazeal 2006). Assuming hardcoded routines for perception and motion control, as readily available in existing middleware packages (such as Player/Stage and Microsoft Robotics Studio) and perception libraries (such as Lowe's SIFT object recognition package and augmented reality tag-tracking libraries, such as ARTag), we focus on learning decision-making policies $\pi : \hat{s} \rightarrow a$ that map perceived robot state $\hat{s}$ to robot actions a.

We claim, at this point in time, that facilitating long-term data collection and human guidance is the primary challenge, more so than algorithm development, for lifelong robot learning. Many approaches to robot policy development have been pursued, ranging from primarily hand-coded domain-specific algorithms (Edsinger 2007) to general nondeterministic adaptive methods, such as learning with Partially Observable Markov Decision Processes (Kaelbling, Littman, & Cassandra 1998). While this space of algorithms and methods is well covered, less attention has been paid to human-robot interfaces that will facilitate the longitudinal human-robot interaction and guidance to enable tractable lifelong robot learning. Notable work in this area includes experiments on managing robot teams (Goodrich *et al.* 2007) and interface design (Sawaragi & Horiguchi 2000). Towards this end, we discuss our experiences using the Nintendo Wiimote as one possible human-robot interface. Through various public demonstrations of our implemented system, we have informally observed that people find Wiimote interfaces "fun" and "engaging", usable within the space of a few minutes, applicable to various robot platforms (e.g., iRobot Create, Sony Aibo, DLR robot hand), and easier to control than standard gamepads/joysticks.

Nintendo Wii remote

My summer project began by attempting to establish control of a robot via a Nintendo Wii remote. Being less than a year old at the time of writing, the Wii remote was a relatively new technology. It was also an exciting piece of technology, because it represented a new way of playing video games, as advertised by Nintendo.

Released in December 2006, the Nintendo Wii Remote (or Wiimote, shown in Fig. 1) is an inertial control interface for video games that is fundamentally different from traditional gamepad/keyboard/mouse devices. The primary innovation of the Wiimote is its ability to localize itself within 2 rotational and 3 translational degrees of freedom. This localization is performed with a reasonable degree of accuracy which is well complemented by the Wiimote's economical feasibility and compelling aesthetic. Rotational localization occurs with the help of three inertial sensors (accelerometers/gravimeters) that measure the direction of gravity along roll, pitch, and yaw axes. Translational localization is performed through triangulation against infrared light (IR) emitted by an external "sensor bar". The IR is sensed by the Wiimote through a built-in IR-sensitive chip. In addition, a Wiimote can receive input from 12 traditional gamepad buttons that can be used in complement with its localization.

Figure 2: A boy controlling a SmURV at RoboBusiness '07

The Wiimote communicates with other devices using the Bluetooth wireless communication. There are certain events that cause the Wiimote to send a packet of updated state information to its connected device. Those events include button presses, button releases, changes in the data from the accelerometers or the IR sensor, and changes in Wiimote extension devices. The Nunchuk (shown on the left in Fig. 1) is one such extension. It physically connects to the Wiimote and adds a second set of 3 accelerometers along with 2 trigger-style buttons and an analog joystick. This combined Wiimote/Nunchuck interface allows for two-handed user input. For robot control, this has made the parallel coordination of navigation and manipulation easier in our informal experience.

Robot control via wiimote

In the same way that Nintendo was trying to reach out and get a new audience to play games with its easy, intuitive controller, I worked during the summer on developing a new interface for robot control that anybody, of any age, would be able to use. Because the Wiimote was so new, very few people had done much with it and it was difficult to learn how to use, but eventually I was able to turn raw data from the Wiimote into parameters that could control a robot.

iRobot Create/Roomba control

Our initial work into Wiimote-based robot control began using a single Wiimote to control our Small Universal Robotic Vehicle (SmURV), pictured in Fig. 2. The SmURV platform uses an iRobot Create as a mobility base for a 1.2 GHz Mini-ITX computer. The total cost of a SmURV (including an iRobot Create, computer parts, and a firewire camera) is $750 USD. In terms of software, the SmURV runs a stripped-down Linux distribution from a flash memory card and is controlled directly by the Player robot middleware (Gerkey, Vaughan, & Howard 2003) through a serial interface. We wrote a simple Player client that talks to the Wiimote through a Bluetooth USB dongle and converts Wiimote events into Player commands for the robot. The Wiimote interface for the SmURV is quite simple: the robot is engaged by holding down the trigger on the bottom of the Wiimote, and, once engaged, the robots forward/backward and rotational velocities are controlled by tilting the Wiimote up/down along its pitch axis and twisting left/right along its roll axis, respectively.

Sony Aibo control

In addition to the SmURV, the other robot platform that we primarily worked with was the Sony Aibo. We chose the Aibo for several reasons. First, we already had done a significant amount of development for the Aibo through our team for Robocup, and we wanted to continue to use this platform in our current work. Second, Aibos are very popular among researchers already and are a common sight at many conferences. Third, the Aibo is a more entertaining and fun-to-use platform than the SmURV. It looks a lot friendlier (dogs are currently cuter than Roombas), and instead of only being able to turn and drive forwards or backwards, the Aibo can also move sideways and execute a large variety of moves that are useful in soccer, including kicks, blocks, and a few other tricks as well.

Figure 3: Girl at AAAI '07 in Vancouver controlling an Aibo with our Wiimote and Nunchuck control system

One of the challenges we faced in using the Wiimote to control a soccer-playing Sony Aibo robot dog was how to map a small amount of user input into the large range of control parameters provided by the Aibo's 18 degrees of freedom (DOFs). Further, the Aibo must also be coordinated to perform both navigation/locomotion and manipulation functions in the course of playing soccer. We thus simplify the control problem by way of a hand-coded motion controller, producing three real-valued DOFs in the head (along with a binary-valued jaw), three real-valued walking gait DOFs (forward, sideways, rotation), and various discrete-valued "soccer moves" (kicks, blocks, ball trapping and control) that involve ball manipulation in some form.

Previously, we had attempted using dual-analog gamepads to control these degrees of freedom. These gamepads consisted of two analog joysticks, a discrete directional pad, and various buttons. Following the model of standard first-person video games, robot locomotion and head movements were controlled by analog joysticks. Each soccer move (kick,block, and trap) was associated with an individual button to trigger their execution.

This control interface, however, seemed to have a steep learning curve that often did not lead to usable human control of the Aibo. We attribute this circumstance to several factors, two of which we address. First, the analog sticks were often used in a "full-throttle" manner, effectively having little more benefit than the discrete directional pad. The real world nature of the robotics domain appears to amplify this phenomenon which we attribute to user impatience to observe the effect of their actions. Second, we believe our users found it difficult to navigate when both walking and moving the head are each controlled with a single thumb.

Through the Wiimote/Nunchuck, we are able to provide an interface that (informally stated) is more usable for Aibo soccer. In our interface, we decided to separate locomotion and manipulation-related functions by having each be controlled entirely by a different hand of the user. Because the head is the Aibo's primary manipulator, the Wiimote is used exclusively to control the head and soccer moves of the robot. The orientation and directional pad of the Wiimote controlled the robot's head and soccer moves, respectively (shown in Fig. 8). On the directional pad, up was mapped to "kicking" with the chest (shown in Fig. 9), left/right mapped to left/right lunge blocks, and down executed a block with both forward limbs (shown in Fig. 7). Additionally, the Wiimote's "A" button was used to execute a trapping motion for acquiring the ball between the robot's chin and chest and the floor (shown in Fig. 6).

Locomotion by forward/backward/turning motions was controlled by the Nunchuck in a similar manner to that of our SmURV controller. In addition, the analog stick is used for sideways strafing. This interface not only separates locomotion onto the Nunchuck and manipulation onto the Wiimote, but also separates head control (performed mostly by the wrist) and instantiation of soccer moves (performed mostly by the thumb). Further, Wiimote control relies only on a user's proprioception about two wrist and one finger DOF rather than the more exteroceptive sensing of a joystick with respect to its base.

The point about engaging the interest and human-robot interaction with kids through Wiimote control becomes magnified for the Aibo. For example, at AAAI, two adolescent boys along with a young girl (shown in Fig. 4) spent over 30 minutes controlling robots around the venue and chasing a ball.

Results

After establishing the connection between Wiimote and robot, I worked on the controls themselves, tightening the responsiveness so that even a small movement of the wrist would be reflected in the Aibo's motions, changing the button-mappings to be more intuitive, and improving general ease of use. I even found some beta-testers to test the control scheme and give feedback. The end result was a simple, intuitive control interface where the user need only twist their wrist to make their Aibo play soccer in the manner the user intended.

Our first demo, at Robo-Business 2007 in Boston, was not run by us, but rather part of iRobot's demos of what can be accomplished with their Create. However, it was still a great opportunity for us to see how the greater public responded to using a Wiimote to control a robot (in this case, only a SmURV, we had not yet begun the work with the Aibo). We observed the following: First, kids seem to love Wiimote robot control. They understand the controls quickly and are then able to channel their general interest in robots into actual behavior. Second, humans appear to have a strong preference for having a direct line of sight to the robot. Our attempts to provide the user only with the robot's onboard camera and sensing by running the Player client remotely met with limited user satisfaction, probably due to camera latency and limited field of view. Third, the SmURV's low cost, combined with its ruggedness, allows human operators to be rough with the robot platform, enabling a more natural human-robot dynamic without the interference of overly cautious researchers.

Figure 4: Both small children and adults alike are enthralled by Aibos being controlled via Wiimote.

Our second demo, at Robocup 2007 in Atlanta, GA, was very different from that of Robo-Business. Now, instead of demoing the SmURV, we were using the Aibo, and instead of presenting to members of the robotics business community, we were presenting to a tighter-knit community of programmers who were all intimately familiar with the Aibo and its capabilities. The highlight of our demo was an unofficial game between 4 of our Aibos, all controlled remotely by Wiimote, versus an actual Robocup team, running autonomously. People seemed to find the interface very engaging, and many of them seemed to pick up the controls very quickly. Again, we noticed that children seem to be particularly attracted to Wiimote control of robots. This could be because of the growing popularity of the Wii, but it could also be due to the fact that we had succeeded in making an intuitive interface.

Our third demo occurred at AAAI 2007 in Vancouver. The response was much the same as in Atlanta, with visitors of all ages, genders, and nationalities trying the interface and commenting on its ease-of-use. In summary, we got a lot of

Starting control loop with keyboard.

Figure 5: Preview of the game for gathering data.

Figure 6: Aibo performing a trap.

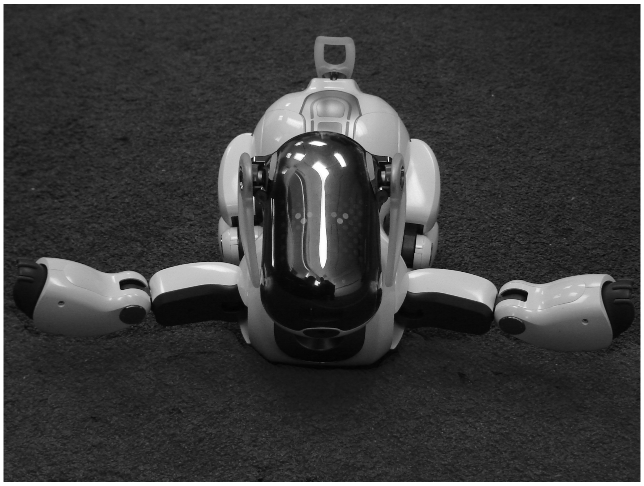

Figure 7: Aibo performing a block.

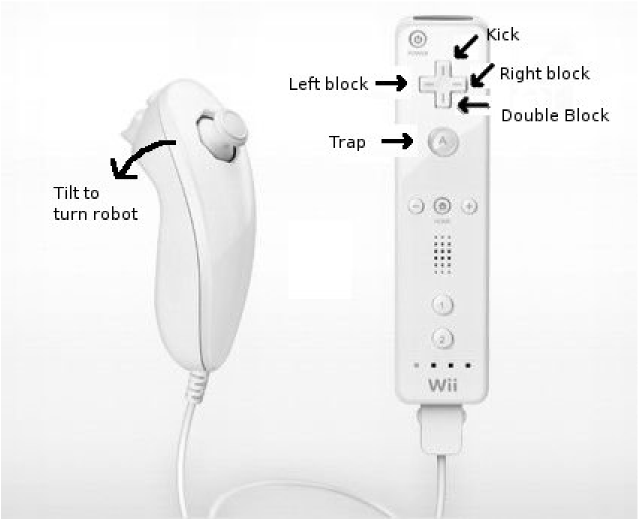

Figure 8: The control mappings.

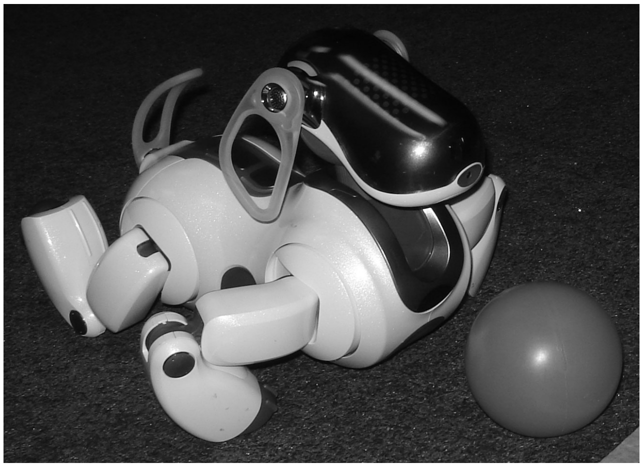

Figure 9: Aibo kicking a ball.

positive feedback on the system, including suggestions of additional button-action mappings and other small improvements.

General observations

We have found informally, through the public demonstrations mentioned above and at local venues, that our control system involving the Wiimote and Nunchuck is very intuitive and easy to use. With only about 15-30 seconds of explanation, we have seen users of all ages ranging from 3 to 75 pick up these controllers and within a minute be able to fully control an Aibo. The simple act of twisting your hand to control all the walking and head motions is very simple and quickly becomes reflexive, allowing the users conscious to focus on meeting task-oriented objectives rather than fighting to get an awkward control system to do what they want. More often than not, after "driving" an Aibo around for even a few seconds, users will comment on the facility and intuitive nature of the control system.

Future work

Although the summer has ended, there is still plenty of work to be done. The true purpose of the Wiimote control system was to develop an intuitive method of control to make it easier for people to demonstrate their desired control policy. The next step is to take the data from those users controlling robots with Wiimotes and learn from it.

Our plan for the development of a lifelong learning robotic system is intimately connected to the Wiimote control system. The kind of learning we are interested in is learning from demonstration (Grollman & Jenkins 2007), where a human controller demonstrates to the learning program how to perform various activities, in this case playing soccer. However, the algorithms in this space require a great deal of data in order to create good policies. Therefore, we are in the midst of developing a free downloadable game where users can control a soccer-playing robot from the comfort of their homes, over the Internet. Users will be able to log on and play games, using either the Wiimote interface or any of the other supported interfaces (such as the standard keyboard and mouse) to help us gather data to learn from.

To further this goal, we are working on various ways to improve a user's ability to provide a clean, successful control policy to learn from. One such area is in the actual system of control itself, hence the initial development of the intuitive and easy-to-use Wiimote interface. Another such area is presentation of the robot's perceived state to the user. As stated in the Results section, we observed some frustration from users while trying to control a SmURV or Aibo without a direct line of sight to the robot (i.e. seeing just the camera view instead). This presents a future problem for learning given that the human and the robot will have completely different notions of state. We envision that a combination of augmented reality techniques along with a presentation of robot perception, as commonly used in networked video games, could provide a working solution in the future, and this represents another area of research that we are exploring.

Conclusion

There is still a long way to go before fully autonomous learning robots are available. However, the work that I have done towards such a goal has been an amazing experience for me, and it continues. When I first took this job in the summer, I had no idea where it would take me, but take me away it did.

In our goal to take the development of robot control policies out of the hands of programmers and make it available to the consumer, we have pinpointed a field of AI where undergraduates can very easily and quickly become involved: human-robot interactions (HRI). Specifically, the use of novel, exciting controllers, such as the Nintendo Wiimote, is a great method for teaching undergraduates about various aspects of AI, including control loops, machine learning, robots, and more. It is a field where a lot of technical knowledge isn't necessary to get involved, and successes can happen quickly, keeping the timeline of a project within the sights of undergraduate students. Furthermore, lifelong robot learning aims to bring many people into the fold of robotics, not just through programming or computer science, but instead by enabling them to create robot control policies without any knowledge of programming or computer science. We are hoping that more educators will notice the field of HRI as an easy way to quickly get undergraduate students involved in computer science, artificial intelligence, and robotics.

References

Edsinger, A. L. 2007. *Robot manipulation in human environments*. Ph.D. Dissertation, Cambridge, MA, USA. Adviser-Rodney A. Brooks.

Gerkey, B.; Vaughan, R.; and Howard, A. 2003. The player/stage project: Tools for multi-robot and distributed sensor systems.

Goodrich, M. A.; McLain, T. W.; Anderson, J. D.; Sun, J.; and Crandall, J. W. 2007. Managing autonomy in robot teams: observations from four experiments. In *HRI '07: Proceeding of the ACM/IEEE international conference on Human-robot interaction*, 25–32. New York, NY, USA: ACM.

Grollman, D. H., and Jenkins, O. C. 2007. Learning robot soccer skills from demonstration. In *6th IEEE International Conference on Development and Learning*.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101(1-2):99–134.

Sawaragi, T., and Horiguchi, Y. 2000. Ecological interface enabling human-embodied cognition in mobile robot teleoperation. *Intelligence* 11(3):20–32.

Thomaz, A. L., and Breazeal, C. 2006. Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance.