2-3 Using Humanoid Robots to Study Human Communication

Christopher G. Atkeson
ATR Information Sciences Division

The ATR ISD CyberHuman Project uses humanoid robots to study human behavior and communication. A current focus is learning from demonstration, where a person communicates a skill by showing it to another person or machine. Machine perception of human movement, translating actions and goals, and learning from practice are important ingredients in our approach to learning from demonstration.

Keywords
Humanoid Robots, Learning from Demonstration, Learning from Practice, Machine Vision

The goal of the ATR ISD CyberHuman Project is to develop computational models of human behavior, so that we can support human-human and human-machine communication more effectively. In this article, I will describe some of our work on perception and generation of human motion. We use humanoid robots to test our theories and behavioral algorithms. Using a humanoid robot as a research tool forces us to deal with a complex physical apparatus and complex tasks. Our work excites a lot of public interest, but we have to meet high standards, because observers expect human level competence from a machine with a human form. Humanoid robots have tremendous potential in society, both to serve humans directly and to operate in spaces designed for humans. We envision a future in which humans can explore remote and/or dangerous experiences by communicating with a robot. We expect it will be easier for humans to interact with and control robots with human form. We have an opportunity to develop ways to make it easier to program behavior in a humanoid robot, and potentially in other machines and computer systems as well, based on how we program behavior in our fellow humans.

We will describe our work with our current humanoid robot DB (www.erato.atr.co.jp/DB/), a hydraulic anthropomorphic robot with legs, arms, a jointed torso, and a head (Fig. 1). Several projects are using this robot as a test bed, including the Kawato Dynamic Brain Project, an Exploratory Research For Advanced Technology (ERATO) project funded by the Japan Science and Technology Agency. This robot is unique in the world due to its human-like form and mechanical capabilities. Many international researchers come to ATR to work with this robot.

We have already demonstrated several simple behaviors, including juggling a single ball by paddling it on a racket, learning a folk dance by observing a human perform it[4], robot drumming synchronized to sounds the robot hears (karaoke drumming)[5], juggling 3 balls, performing a Tai Chi exercise in contact with a human[2], and various oculomotor behaviors[6]. We are focusing our research on learning (especially learning from demonstration).

We are interested in how humans and machines can learn from sensory information

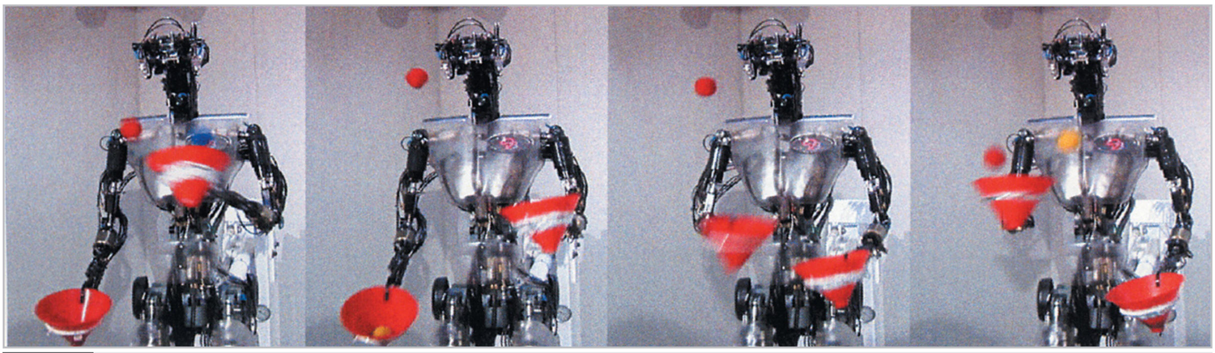

Fig.1 *The humanoid robot juggling 3 balls, using kitchen funnels for hands.*

in order to acquire perceptual and motor skills. For this reason, we are exploring neural networks, statistical learning, and machine learning algorithms. Learning topics that we investigate fall into several areas including learning from demonstration and reinforcement learning.

1 Learning from Demonstration

A major focus of our work with the humanoid robot is learning from demonstration. It typically takes a human a long time to program one of our anthropomorphic robots to do a task. How can we reduce the cost of communicating with and controlling complex systems? One way we instruct our fellow human beings is to show them how to do a task. It is amazing that such a complex sensory input is useful for learning. How does the learner know what is important or irrelevant in the demonstration? How does the learner infer the goals of the performer? How does the learner generalize to different situations? Our hope is that human-like learning from demonstration will greatly reduce the cost of programming complex systems. In addition, we expect humanoid robots to be asked to perform tasks that people do, which typically involve human-like motions which can easily be demonstrated by a human.

We also believe that learning from demonstration will provide one of the most important footholds to understand the information processes of sensori-motor control and learning in the brain. Humans and many animals do not just learn a task from scratch by trial and error. Rather they extract knowledge about how to approach a problem from watching others performing a similar task, and based on what they already know. From the viewpoint of computational neuroscience, learning from demonstration is a highly complex problem that requires mapping a perceived action that is given in an external (world) coordinate frame into a totally different internal frame of reference to activate motor neurons and subsequently muscles. Recent work in behavioral neuroscience has shown that there are specialized neurons (“mirror neurons”) in the frontal cortex of primates that seem to be the interface between perceived movement and generated movement, i.e., these neurons fire very selectively when a particular movement is shown to the primate, but also when the primate itself executes the movement. Brain imaging studies with humans are consistent with these results.

Research on learning from demonstration offers a tremendous potential for future autonomous robots, and also for medical and clinical research. If we can communicate with machines by showing, our interaction with machines would become much more natural. If a machine can understand human movement, it can also be used in rehabilitation as a personal trainer that watches a patient and provides specific new exercises to improve a motor skill. Finally, the insights into biological motor control developed in learning from demonstration can help to build adaptive prosthetic devices that can be taught to improve the performance of a prosthesis.

One working hypothesis is that a perceived movement is mapped onto a finite set of movement primitives that compete for perceived action. Such a process can be formulated in the framework of competitive learning. Each movement primitive predicts the outcome of a perceived movement and tries to adjust its parameters to achieve an even better prediction, until a winner is determined. In preliminary studies with anthropomorphic robots we have demonstrated the feasibility of this approach. Nevertheless, many open problems remain for future research. We are also trying to develop theories on how the cerebellum could be involved in learning movement primitives.

To explore these issues we have implemented learning from demonstration for a number of tasks, ranging from folk dancing to various forms of juggling. We have identified a number of key challenges. The first challenge is to be able to perceive and understand what happens during a demonstration. The second challenge is finding an appropriate way to translate the behavior into something the robot can actually do. Although our current robot is humanoid, it is not a human. It has more restrictive joint movement limits, is weaker, and its maximum speeds are slower than a human. It has many fewer joints and ways to move. A third challenge is that there are many things that are hard or impossible to perceive in a demonstration, such as muscle activations or responses to errors that do not occur in the demonstration. The robot must fill in the missing information using learning from practice. Solving these challenges is greatly facilitated by having the robot be able to perceive the teacher's goal.

2 Perceiving Human Movement

In order to understand a demonstration of a task, the robot must be able to see what is happening. We have focused on the perception of human movement. We are exploiting our knowledge of how humans generate motion to inform our perception algorithms. For example, one theory of human movement is that we move in such a way as to minimize how fast muscle forces change[3]. This theory about movement generation can be used to select the most likely interpretation of ambiguous sensory input[7].

Our first thought was to borrow motion capture techniques from the movie and video game industry. However, we found that the requirements to actually control a physical device such as the humanoid robot, rather than draw a picture, required substantial modifications of these techniques. We have experimented with optical systems that track markers, systems where the teacher strapped on measurement devices, and vision-based systems with no special markers.

The organizing principle for our perception algorithms is that they should be able to recreate or predict the measured images based on the recovered information. In addition, the movement recovery is made more reliable by adding what are known as "regularization" terms to be minimized. These terms help resolve ambiguities in the sensor data. For example, one regularization term penalizes high rates of estimated muscle force change. We also process a large time range of inputs simultaneously rather than processing images or measurements taken at a single time, so we can apply regularization operators across time as well and easily handle occlusion and noise. Thus, perception becomes an optimization process which tries to find the underlying movement or "motor program" that predicts the measured data and deviates from what we know about human movement the least.

In order to deal with systems as complex as the human body and the humanoid robot, we had to use a representation with adaptive resolution. We chose B-spline wavelets. Wavelets are removed when their coefficients are small, and added when where there is large prediction error. We have also developed large scale optimization techniques that handle the sparse representations we typically find in the observed data. These optimization techniques are also designed to be reliable and robust, using second order optimization with trust

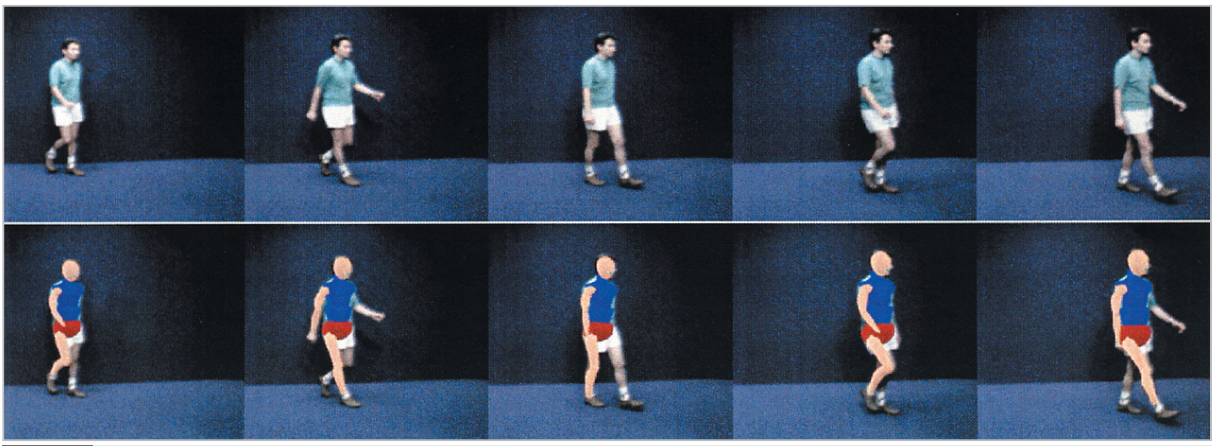

Fig.2 *Perceiving human motion. The top row of frames show a human walking, and the bottom row of frames show how well our perception is tracking the motion by overlaying a graphical model where the perception system believes the human body parts to be.*

regions and also using ideas from robust statistics allowing us to take into account only the relevant data while ignoring background information and noise which should not influence the interpretation of the perceived actions. Fig. 2 shows an example of our perception algorithms applied to frames from a high speed video camera.

3 Translating Movement and Inferring Goals

We used an Okinawan folk dance "Kachashi" as one test case for learning from demonstration[4]. We captured movements of a skilled performer. After using the perception techniques described above, we found that the motions of the teacher exceeded the joint movements the robot was capable of. We had to find a way to modify the demonstration to preserve the "dance" but make it possible for the robot to do. We considered several options:

(1) Scale and translate the joint trajectories to make them fit within robot joint limits. The Cartesian location of the limbs is not taken into account.

(2) Adjust the visual features the robot is trying to match until they are all within reach. This can be done by translating or scaling the images or three dimensional target locations. It is not clear how to do this in a principled way, and the effects on joint motion are not taken into account.

(3) Build the joint limits into a special version of the perception algorithms, so that the robot can only "see" feasible postures in interpreting or reasoning about the demonstration. This approach trades off joint errors and Cartesian target errors in a straightforward way.

(4) Parameterize the performance in some way (knot point locations for splines, for example) and adjust the parameters so that joint limits are not violated. Human observers score how well the "style" or "essence" of the original performance is preserved, and select the optimal set of parameters. This is very time consuming to do, unless it is possible to develop an automatic criterion function for scoring the motion.

We implemented the first option. It is clear that we should also consider the alternative approaches. We learned from this work that we need to develop algorithms that identify what is important to preserve in learning from a demonstration, and what is irrelevant or less important. For example, we have begun to implement catching based on learning from demonstration (Fig. 3), where the learned movement must be adapted to new requirements, such as the ball trajectory[4]. For catching what is important is that the hand intercept the ball at the right place and time in space, and the joint angle trajectories are secondary.

We have begun to implement learning how

to juggle three balls from demonstration on the humanoid robot. We have found that in this case actuator dynamics and constraints play a crucial role. Because the hydraulic actuators limit the joint velocities to values below that observed in human juggling, the robot needs to significantly modify the observed movements in order to juggle successfully. We have manually implemented several feasible juggling patterns, and one pattern is shown in Figure 1. Something more abstract than motion trajectories needs to be transferred in learning from demonstration. The robot needs to be able to perceive the teacher's goals to perform the necessary abstraction. We are currently exploring alternative ways to do this.

4 Learning from Practice

After the robot has observed the teacher's demonstration, it still must practice the task, both to improve its performance and to estimate quantities not easily observable in the demonstration. In our approach to learning from demonstration the robot learns a reward function from the demonstration, which then allows it to learn from practice without further demonstrations[1]. The learned reward function rewards robot actions that look like the observed demonstration. This is a very simple reward function, and does not capture the true goals of actions, but works well for many tasks. The robot also learns models of the task from the demonstration and from its repeated attempts to perform the task. Knowledge of the reward function and the task models allows the robot to compute an appropriate control mechanism.

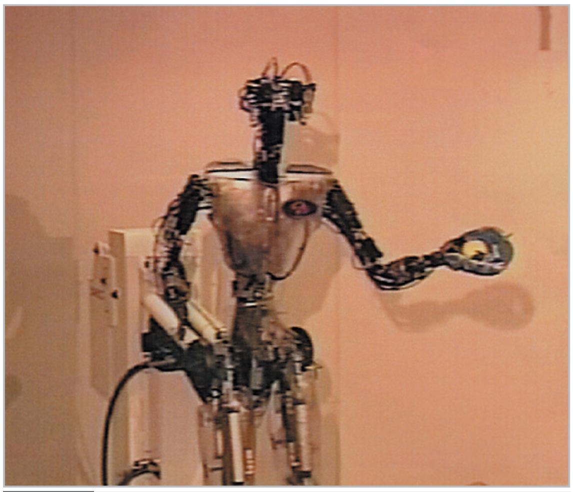

Fig.3 A frame of motion showing the end of a catching sequence.

Lessons learned from implementations of learning from practice include:

- Simply mimicking demonstrated motions is often not adequate.
- Given the differences between the human teacher and the robot learner and the small number of demonstrations, learning the teacher’s policy (what the teacher does in every possible situation) often cannot be done either.
- However, a task planner can use a learned model and reward function to compute an appropriate policy.
- This model-based planning process supports rapid learning.
- Both parametric and nonparametric models can be learned and used.
- Incorporating a task level direct learning component, which is non-model-based, in addition to the model-based planner, is useful in compensating for structural modeling errors and slow model learning.

5 Future Goals

Future goals of the CyberHuman Project include communicating style, more complete behaviors, and interacting with systems that have a continuous existence. For example, many video game players would like to create characters that behave like the human player, including aspects of their personality. Currently we can communicate isolated skills, but we look forward to communicating much more complete models of behavior.

References

1 C. G. Atkeson and S. Schaal, "How can a robot learn from watching a human?", In Proceedings of the Fourteenth International Conference on Machine Learning (ICML'97), pages 12-20. Morgan Kaufmann, San Francisco, CA, 1997.

2 J. G. Hale and F. E. Pollick, "Sticky hands' interation with an anthropomorphic robot", In 2000 Workshop on Interactive Robotics and Entertainment (WIRE-2000), 2000.

3 M. Kawato, "Internal models for motor control and trajectory planning", Current Opinion in Neurobiology, 9:718-727, 1999.

4 M. Riley, A. Ude, and C. G. Atkeson, "Methods for motion generation and interaction with a humanoid robot: Case studies of dancing and catching", In 2000 Workshop on Interactive Robotics and Entertainment (WIRE-2000), 2000.

5 S. Schaal, S. Kotosaka, and D. Sternad, "Nonlinear dynamical systems as movement primitives", In IEEE International Conference on Computational Intelligence in Robotics and Automation (CIRA'99), 1999.

6 T. Shibata and S. Schaal, "Fast learning of biomimetic oculomotor control with nonparametric regression networks", In IEEE International Conference on Robotics and Automation (ICRA'00), 2000.

7 A. Ude, C. G. Atkeson, and M. Riley, "Planning of joint trajectories for humanoid robots using B-spline wavelets", In IEEE International Conference on Robotics and Automation (ICRA'00), 2000.

***Christopher G. Atkeson**, PhD*

Invited Researcher, Cyberhuman Project Information Sciences Division ATR International / Associate Professor, Robotics Institute Carnegie Mellon University

Machine Learning