Exploiting Human Motor Skills for Training Bipedal Robots

Undergraduate Honors Thesis

Adam Setapen
Department of Computer Sciences
University of Texas at Austin
asetapen@cs.utexas.edu

Supervising Professor: Peter Stone

April 20, 2009

Abstract

Although machine learning, reinforcement learning, and learning from demonstration have improved the rate and accuracy at which robots can gain intelligence from humans, they haven't reached the rapid rate at which humans are able to acquire new knowledge. Many systems that exploit imitation learning use simple positive and negative reinforcement, and place the burden of learning completely on the computer. This neglects the expressive capabilities of humans, as well as their remarkable ability to quickly refine motor skills. While passive dynamics offers the most human-like locomotion for bipedal robots, it also relies on particular design specifications. This thesis presents a general Framework for Interactive Control of a Humanoid by Motion Capture (*FICHMC*), that offers rapid motion development for large classes of bipedal robots. Essentially, a human in a motion-capture laboratory "puppets" a biped, with a real-time mapping from human to robot. The training process requires no technical knowledge and provides a natural interface for humans to directly transfer skills to robots.

Contents

1 Introduction **3**

2 Background **5**
2.1 Human-robotic interaction . . . 5

3 Hardware systems **6**
3.1 PhaseSpace IMPULSE Motion Capture System . . . 6
3.2 The Aldebaran Nao . . . 6
3.2.1 Kinematic Model . . . 7

4 *FICHMC*: Framework for Interactive Control of Humanoids through Motion Capture **9**
4.1 Justifying our approach . . . 9
4.2 Implementation details . . . 10
4.3 User Interface and Training . . . 13

5 Results **14**
5.1 Car-Park Task . . . 15

6 Discussion **16**
6.1 Bipedal locomotion . . . 16
6.1.1 Passive dynamic walking . . . 17
6.1.2 Extending the problem . . . 17
6.1.3 Alternative approaches . . . 18
6.1.4 Passive dynamics on the Nao . . . 18
6.2 Related Work . . . 19
6.3 Future Work . . . 19

7 Conclusion **21**

1 Introduction

The past few decades have enabled computers to surpass humans at many jobs, specifically those that are computationally demanding. Hardware is getting faster (and cheaper), the internet contains a massive amount of useful data, and innovative algorithms are able to parse and interpret this data with remarkable speed. However, despite the fact that modern computers are capable of outperforming humans at many tasks, there is still a large group of jobs for which human skill surpasses computers. For example, humans are easily able to walk across a room, identify a specific friend, and give them a handshake - this is no trivial task for a robot. However, if you ask a human to determine the average population of every city in Pennsylvania without using a computer, it would take them days, if not weeks. Computers are now completely immersed in many people's lives, and recent breakthroughs in robotics and artificial intelligence suggest that in the future, robots will also play an integral part in our everyday routine.

Machine learning, the study of algorithms that improve automatically through experience, has drastically improved the rate at which robots can learn. Recently, machine learning algorithms have seen great success in the training robots to move quickly and efficiently. For example, the work of Saggar, D'Silva, Kohl, and Stone presents a policy gradient machine learning algorithm specialized for finding fast quadruped locomotion while ensuring a stable camera [31] (which is necessary for other essential tasks in robotics such as image segmentation and localization). There are numerous case studies in which machine learning has been used for on-line and off-line performance improvement in multi-agent autonomous robot environments [28].

Human-robotic interaction (HRI) examines the principles and methods by which robots and humans can naturally communicate and interact. In the near future, humanoid robots will not be limited to scientists and researchers, but to the non-technical community as well. As the number of robots we interact with on a day-to-day basis increases, the ability for non-technical users to program, train, and interact with these robots will be vital. Thus, any viable framework for human-robotic interaction should require very little technical knowledge to use. Additionally, HRI systems should aim to make the method of communication between robot and human as natural as possible, namely by providing a convenient interface for the human.

As human-robotic interfaces mature, it is inevitable that agents will need to learn directly from humans. Learning from demonstration (sometimes called imitation learning) is a process in which a robot attempts to learn a task by observing a demonstration, typically performed by a human. Imitation learning is a promising way of transferring knowledge from agent to agent, and work by Dautenhahn and Nehaniv illustrates how many animals use this technique to learn new skills [10]. A good deal of recent work in learning from demonstration uses human feedback as a reward signal to a reinforcement learning or policy-search algorithm [16], [6]. Reinforcement learning, a specialized type of machine learning, encompasses the task of learning what to do - namely how to map states to actions in order to maximize a reward function. The learner is not given any instructions as to which action to take, but must explore the state-space and discern which actions result in the highest reward [34]. Figure 1 shows a typical control-flow diagram of such a framework (specifically from the TAMER framework proposed by Knox and Stone [16]). Note that in this scheme, the learning process is fully contained within the agent, with only its policies influenced by human reinforcement.

While machine learning, reinforcement learning, and learning from demonstration have advanced the speed and accuracy at which robots can learn from humans, they haven't even come close to approaching the rate at which humans are able to acquire new knowledge. In this thesis, we take a new approach through an innovative framework for training biped robots that exploits a system finely tuned

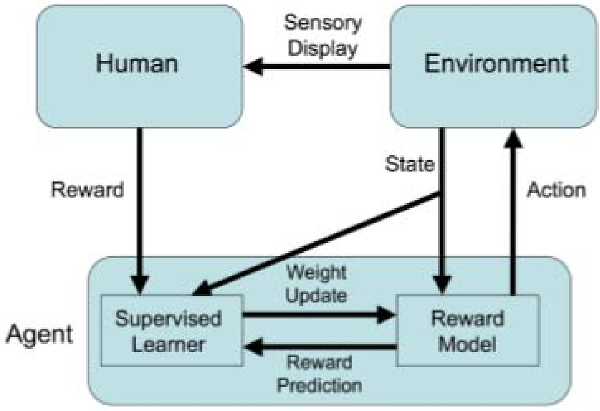

Figure 1: Typical flow for training agents via human feedback[16]

by thousands of years of biological evolution: the human body.

Our work introduces a novel interface between a human in a motion-capture suit and a humanoid robot. In our framework, the learning happens exclusively by the human - not the robot. While the process by which humans are able to learn exceptionally quickly is not yet fully understood, work being done on the neurological basis of learning is steadily shedding light on how humans rapidly acquire and apply new knowledge [19]. Specifically, our framework exploits the ability at which humans are able to learn and refine fine-motor skills. Recent breakthroughs in behavioral motor control have enhanced our understanding of the human brain and illustrate how remarkable our innate capacity for delicate motor control is [32]. Recent work by Muellbacher et al. indicates that given a 60-minute training period, human subjects can rapidly optimize performance of a complex task involving fine motor control [24].

This thesis presents an original framework for training biped robots that has been developed, applied, and tested on a real humanoid robot. Our framework has numerous possibilities for applications - consider a leading brain-surgeon being able to perform operations remotely, or a lazy human teaching a robot how to correctly fold laundry. We will specifically evaluate the rate at which human subjects learn to exploit a direct robot mapping, and one motivation of our approach is to develop an efficient and stable robotic gait.

Due to hardware problems with the legs of our humanoid, we were unable to use our approach to train a walk. However, we have fully implemented a framework for training humanoids by motion capture, and our results indicate that humans are able to quickly improve robotic performance of a task requiring fine-motor skills. We have coined our novel framework *FICHMC*: **F**ramework for **I**nteractive **C**ontrol of **H**umanoids through **M**otion **C**apture. Although our work with *FICHMC* has just begun, we believe that it is a powerful way of harnessing the cognitive flexibility of humans for training large classes of robots.

The structure of this thesis is as follows: Section 2 provides a background of human-robotic interaction. Section 3 describes the motion capture system and humanoid robot that were used. Section 4 presents our motivation, implementation details, and user interface of *FICHMC*, and Section 5 describes some of our results in using *FICHMC*. Section 6 contains a discussion of previous work on bipedal locomotion, presents work related to *FICHMC*, and outlines some possible directions for future work using *FICHMC*. Finally, Section 7 presents our conclusions.

2 Background

This section presents a brief description of current work in human-robotic interaction.

2.1 Human-robotic interaction

Human-robotic interaction is a young and exciting field that investigates the principles required for a natural synergy between humans and robots. Though the field is relatively new, recent work has shed some light on the factors needed to facilitate humanoid robots as cooperative partners for humans. Breazeal et al. argue that "there are many reasons to believe that a social interaction will be the most natural and intuitive way for ordinary people to work with humanoid robots and to teach them" [4]. While there is still a large amount of work being done on finding the perfect ingredients for human-robot interfaces, there have been very successful cases of fully-implemented systems, such as socially assistive robots assisting post-stroke rehabilitation patients [21]. Steinfeld et al. emphasize the need for any viable human-robot interface to cater to non-technical users, and state that when testing any interface it is "critical to recruit subjects having a broad range of knowledge, experience, and expertise" [33]. Although our work doesn't address the societal implications of social robots, there is ongoing work addressing this topic [12].

There are numerous case studies which show that using human feedback to a reinforcement learning or policy-search algorithm significantly improves learning speed [35], [16], [6]. [2]. This illustrates the importance of harnessing human-robot interfaces in order to "design algorithms that support how people want to teach and simultaneously improve the robots learning behavior" [35]. Thomaz and Breazeal coin this paradigm "Socially guided machine learning", where the benefits of machine learning are combined with the intuitive knowledge of humans.

Learning from demonstration (LfD) is one tested way of harnessing the power of a human-robot interface that has shown extremely positive results. However, Argall, Chernova, Veloso, and Browning insist that the LfD community is suffering from the lack of a structured approach. They propose that the problem be broken down into two distinct phases - first gathering examples and then deriving a policy [1].

Breazeal and Scassellati posit that there are four integral questions to consider when designing a system that uses learning from demonstration [5]:

- How does the robot know when to imitate?
- How does the robot know what to imitate?
- How does the robot map observed action into behavior?
- How does the robot evaluate its behavior, correct errors, and recognize when it has achieved its goal?

Our system bypasses the first two questions, as the robot imitates the human in real-time. The robot maps observed actions into behaviors using a deterministic scaling function (that can be tuned by the user). Finally, while our robot doesn't currently evaluate its behavior, we discuss this question in our future work (Section 6.3).

3 Hardware systems

In this section, we describe the motion capture system that was used (Section 3.1) and the humanoid robot that was used (Section 3.2).

3.1 PhaseSpace IMPULSE Motion Capture System

The PhaseSpace IMPULSE motion capture system uses a standard server-client interface, with a flexible API allowing the development of custom clients. The system utilizes 16 high-sensitivity 12.6 megapixel cameras positioned overhead surrounding a 20 by 20 foot tracking area. A human subject wears a black virtual-reality body suit, on which 36 LED markers are strategically placed (7 on each leg and foot, 4 on the right arm, 12 on the left arm and hand, and the rest comprising the body). With a sample rate of up to 480 Hz and a latency of less than 10ms, the PhaseSpace IMPULSE system is a fast and accurate way of capturing even the most subtle human movements. The server receives data from the cameras and disambiguates the LEDs of the body-suit as fixed-point markers. The 36 markers are clustered into groups of four, and each group is updated at 120Hz (thus helping disambiguate unique marker positions). These markers can be combined to represent rigid-bodies - groups of markers that are stationary with respect to one another. Once these markers and rigid bodies are computed by the server, they are made available for reception by a client.

A majority of motion-capture labs also place cameras at the floor of the bounded area, in order to account for a condition in which the line of sight to a particular marker is obstructed by all cameras. However, our lab only has the overhead cameras, which often leads to markers "dropping out" for a brief period. The 480Hz sample rate allows for some flexibility as to when commands are sent to the robot, but there are still times when the confidence of a particular marker is relatively low. Last year, Lee, Kulic, and Nakamura presented a way to recover such "lost markers" using a factorial hidden Markov model [20]. Although dropped markers are not terribly frequent, implementation of this HMM-based estimation could improve data reliability.

3.2 The Aldebaran Nao

We have implemented *FICHMC* on a humanoid robot called the Nao. Developed by Aldebaran Robotics, the Nao is the current platform for the RoboCup Standard Platform League. Utilizing an AMD Geode 500Mhz processor and 256MB of memory, the Nao runs a conventional Linux operating system and includes ethernet and 802.11a/b/g connectivity. Measuring at 23 inches and just under 9.6 pounds, the The Nao has 21 degrees of freedom and body proportions similar to that of a human. Although the version of the Nao we used does not have functioning hands, a version exists with a rotating wrist and actuated fingers. A complete diagram specifying each joint of the Nao can be seen in Figure 2. Each joint houses a dedicated motor - the type in the top portion of the body has a nominal speed and nominal torque of 8810 RPM and 3.84 mNm, while the leg motors have a nominal speed and torque of 6330 RPM and 12.3mNm. Each foot of the robot contains four force-sensitive resistors, with a working range from $0-25N$. The Nao houses an integrated inertial measurement unit with its own processor, which contains a two-axis gyrometer (5% precision and an angular speed of $\sim 500°/sec$) and a three-axis accelerometer (1% precision with an acceleration of $\sim 2G$). Four ultrasound sensors and two 640x480 cameras provide range-sensing and vision capabilities, respectively. A powerful C++ interface, called $NaoQi$, allows for controlling the robot and harnesses an event-based shared memory. Early testing was performed with the Webots simulation software, which accounts for the unique dynamics of the Nao. The UT Austin

Villa Robotic Soccer Team has developed an interface to *NaoQi* using the Lua scripting language to allow for a dynamic programming environment for implementing high-level behaviors.

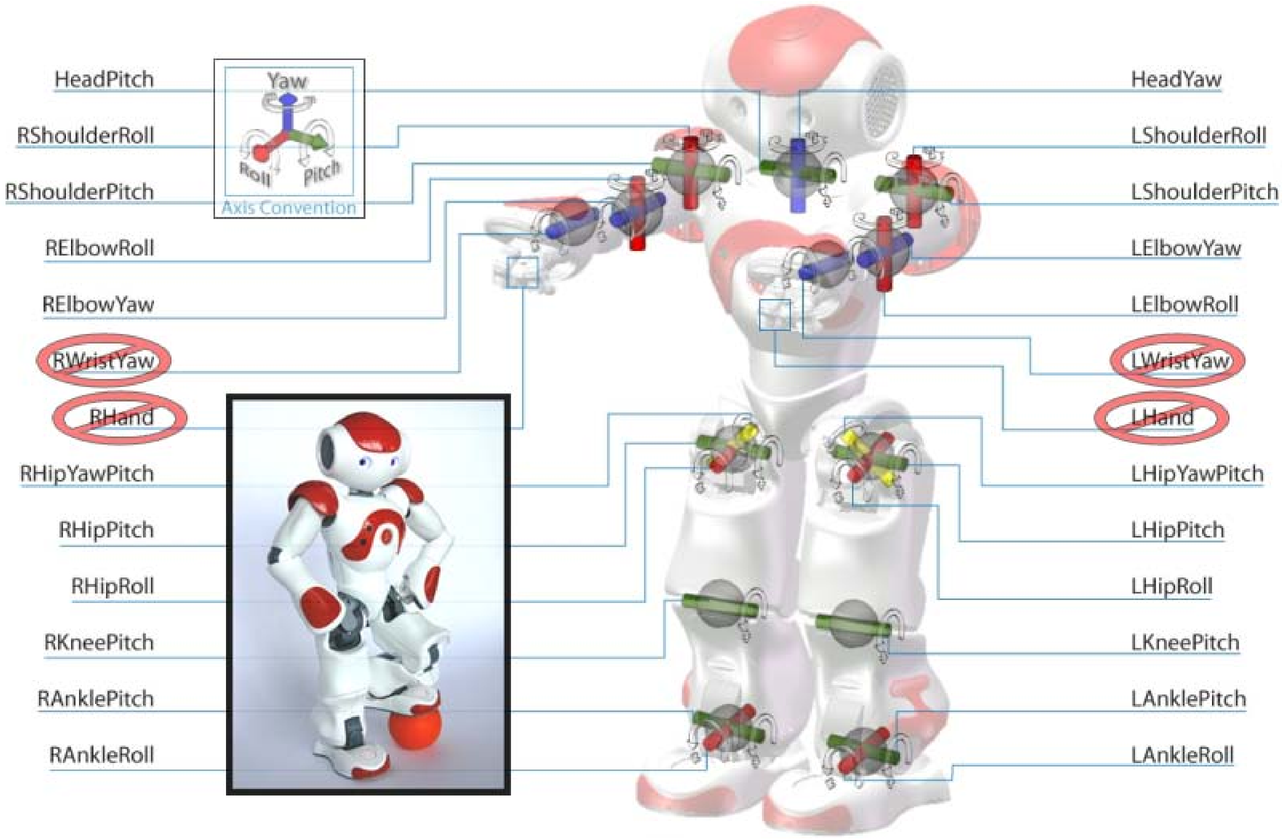

Figure 2: Joint diagram specifying 21 DoF for the Aldebaran Nao

3.2.1 Kinematic Model

The Nao's kinematic model is based on a modified Denavit-Hartenberg representation, a relatively straightforward scheme for describing robotic chain-manipulators. First proposed in 1955 [11], The Denavit-Hardenberg model systematically assigns an orthonormal (x, y, z) coordinate system to every joint in the chain, and has become the standard representation for serial manipulators. Using these individual coordinate systems in each joint, we can simply relate any joint in a chain to the following joint, which results in a complete representation of the Nao's geometry. Figure 3 depicts a simple three-joint chain, using a modified notation with joints connected by links.

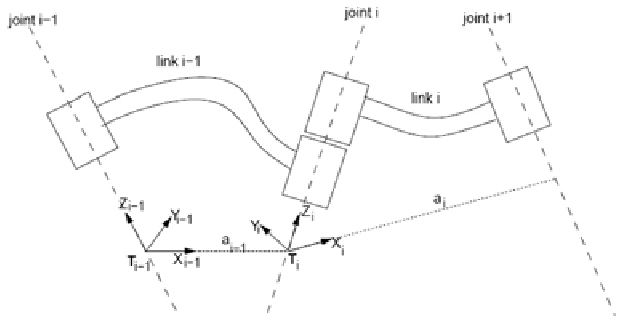

Figure 3: Modified Denavit and Hartenberg scheme for a three-joint chain

Any joint can be related to the following joint using four straightforward parameters, which can be described as follows:

- link length a: offset distance between the Z_{i-1} and Z_i axes along the X_{i-1} axis
- link twist α: angle from the Z_{i-1} axis to the Z_i axis about the X_{i-1} axis
- link offset d: distance from the origin of frame X_{i-1} to the X_i axis along the Z_i axis
- joint angle θ: angle between the X_{i-1} and X_i axes about the Z_i axis

Figure 4 represents the Denavit and Hartenberg schematic of the Nao, comprising a complete model for specifying forward kinematics. Note that the HipYawPitch of the left and right leg are physically bound.

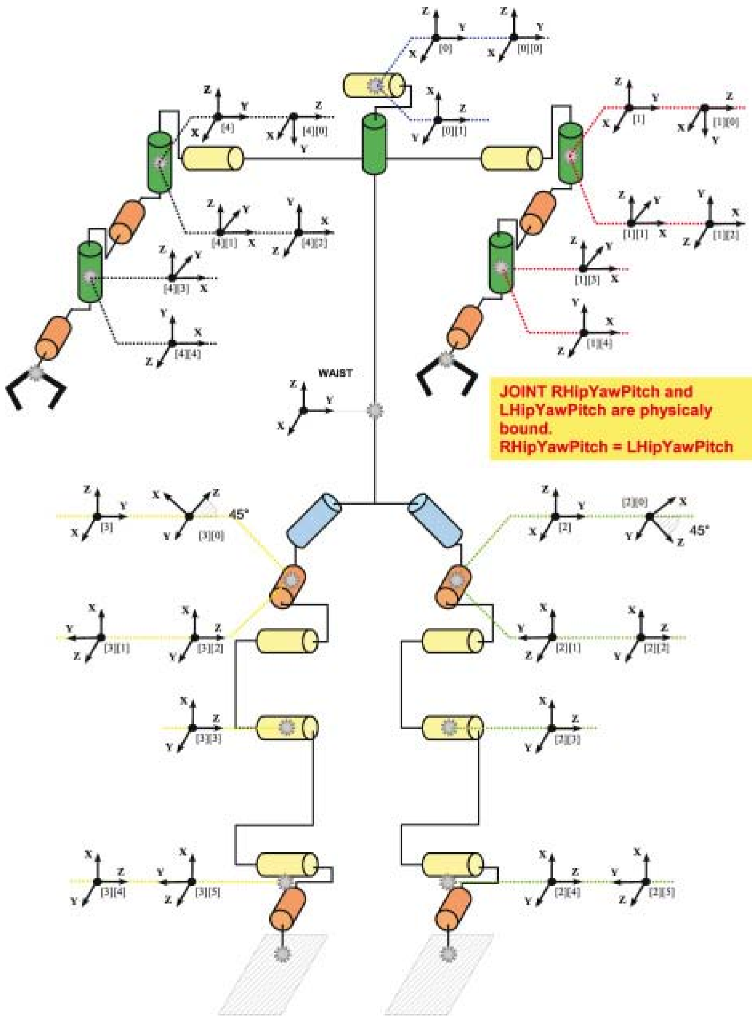

Figure 4: Complete Denavit and Hartenberg Nao Schematic

4 *FICHMC*: Framework for Interactive Control of Humanoids through Motion Capture

The following section describes the *FICHMC* framework. Section 4.1 outlines our motivation for developing *FICHMC*. Section 4.2 describes the implementation details of *FICHMC*, and section 4.3 describes the *FICHMC* user interface and training methodology.

4.1 Justifying our approach

The overall motivation of our approach is that a real-time mapping from a human to a robot will serve as a convenient interface for quickly and systematically training efficient motion sequences. While there is certainly a difference in the natural dynamics of the Nao and Homo sapiens, it is our belief that people's ability to quickly learn fine motor skills can be exploited to improve the gait of the Nao. Essentially, *FICHMC* uses a human to control the robot as a puppet. Even if the mapping from human-sized coordinates to the smaller robot coordinates contains minor errors, humans should be able to learn and exploit these errors. Indeed, modeling the robot's natural dynamics (assuming they are not friction-dominated), might lead to a more power-efficient walk. However, once the mapping is established, we believe that our novel interface will provide a convenient way to quickly develop effective motions for classes of bipedal robots - even robots where physical limitations inhibit passive dynamic walking.

For example, imagine the following situation: the human subject realizes that when training the robot to walk, the robot frequently loses balance and topples forward. The human can try different things to correct this - lengthening of stride, reduction of knee-bend, etc. Humans are so skilled at fine motor control, that minor nuances essential for maximizing performance in bipedal locomotion (typically discoverable through computationally expensive and time-consuming exhaustive methods) may be found with significantly less effort.

One of the motivations for this work is to create a rapid development environment for creating finely-tuned and specialized motions for the UT Austin Villa robotic soccer team. A process that might take weeks using a standard machine learning approach could be approximated in a matter of hours. This would enable Austin Villa to create a large set of tailor-made motions, such as walking up to a ball with proper positioning, kicking a ball, or strafing around an opponent.

4.2 Implementation details

FICHMC has been implemented with two distinct modules - a generalized C++ framework that includes a custom IMPULSE client to connect to the motion-capture server and a Lua behavioral scheme that is run directly on the robot. We have developed a tunable and customizable interface for *FICHMC* that facilitates in coordinating gestures between the motion capture server and the Nao. A fully-functional GUI (written using the cross-platform QT framework) utilizing OpenGL allows viewing of both humanoid and robotic kinematics, supports real-time tuning of mapping scalars, and provides a mechanism for recording motion sequences for later use. The *FICHMC* interface includes a custom IMPULSE client that directly communicates with the motion capture server. The sampling rate of the IMPULSE server is 480Hz, but commands are sent to the robot at 16 frames-per-second due to the frame-rate at which the robot can operate.

The control flow of our interface can be seen in Figure 5. A simple initialization procedure correlates the bounds of the human subject to the bounds of the robot. Initialization is essential for every unique subject that puts on the virtual-reality suit due to the subtleties in different human's body geometries. After initialization is complete, a configuration file can be saved in order to bypass this initial step in further training sessions.

Once initialized, the client captures the markers from the motion capture server (via TCP) and decodes each point to a body part based on unique marker IDs. The decoded packet is then transformed from the coordinate system of the IMPULSE system to that of the Nao. The motion-capture geometry is absolute based at one corner of the moveable grid, whereas the Nao coordinate system places the origin at the robot's center of mass (the stomach-region), with a forward-facing x-axis, a vertical z-axis, and a y-axis pointing towards the left side of the robot. The origin is set at the human's stomach, and a forward-facing x-axis is computed by taking the vector orthogonal to the plane represented by vectors from the stomach to the left and right shoulders. Vectors from the stomach are computed for each motion marker, and these vectors are finally rotated around the new appropriate z-axis. These vectors, now in the robot's coordinate system, are scaled down to the robot's size by considering the subject's initialization information in conjunction with the robot's physical bounds and the interactive sliders of the graphical user interface.

We now have a complete representation of the human's body that is scaled down to the robot's

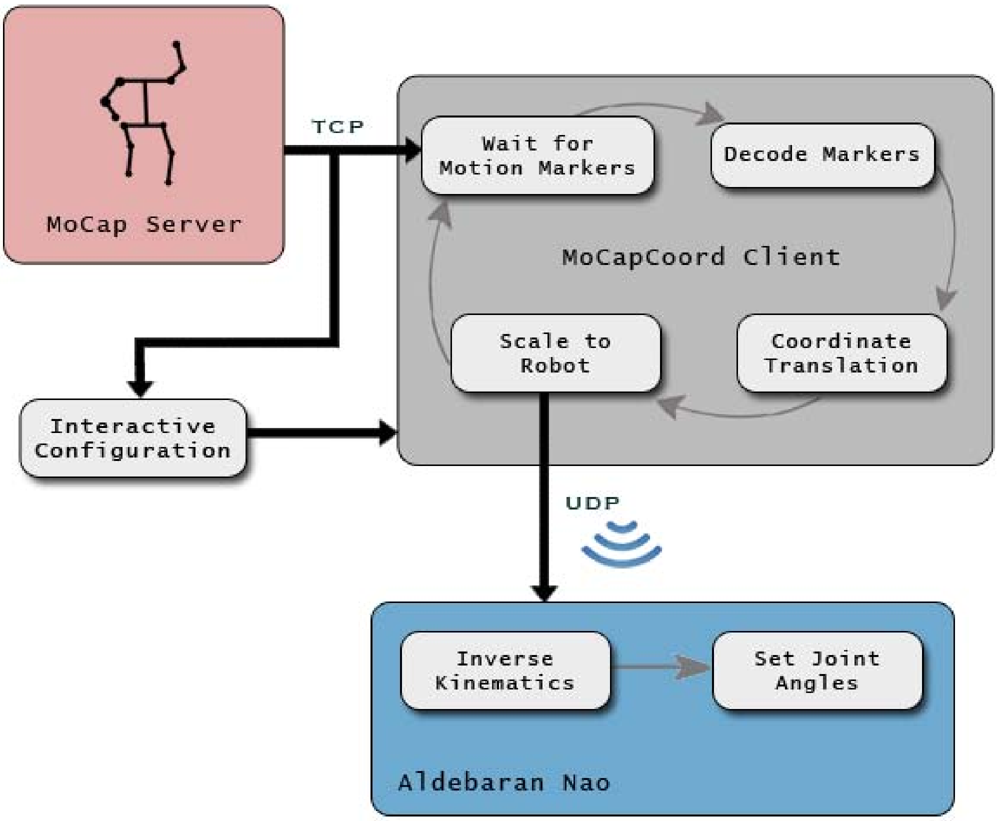

Figure 5: Control flow of the *FICHMC* Interface

coordinate system. We first calculate the end-effector positions of the robots arms and legs. However, bipedal stability is not solely dependent on foot positioning, but also precise ankle orientation. Therefore, we use LEDs on the foot of the subject to determine appropriate angles for the robot's ankle pitch and ankle roll. We package the end-effector positions, the ankle pitch and roll for each foot, and a timestamp into a packet and send this information to the robot via UDP. Pseudocode of the client's algorithm can be seen in Algorithm 1. When the robot receives a packet, it places the corresponding information into its shared memory, which is detected in the robot's Lua module. Using a standard Jacobian approach to inverse kinematics (described in section 3.2.1), the robot sets the joints of each chain in order to set the position of its end-effectors. Finally, the robot sets its AnklePitch and AnkleRoll for each foot, and then waits for more packets. The robot uses a 3rd-degree polynomial smoothing function to interpolate from one motion to the next, resulting in seemingly continuous motion. The robot's motion module is outlined in Algorithm 2.

```
Input: n (number of markers), config (Body Configuration), framerate (Robot framerate)
timestamp = 0;
while ++timestamp do
    n = getMarkers();
    if n == 0 then
        continue;
    end
    MoCapPacket pose;
    foreach marker m do
        decodeToBodyPart(m);
        pose.add(m.x, m.y, m.z);
    end
    if handGestureDetected(LOOP) then
        loopState = nextLoopState();
    end
    // whereframerate = actual framerate / 480Hz
    if timestamp % framerate == 0 then
        if loopState == RECORDING then Recording a loop
            loop.add(pose);
        end
        else if loopState == LOOPING then Playing back a loop
            pose = loop.nextPose();
        end
        Vector forward = orthogonal(pose.lShoulder - pose.stomach, pose.rShoulder - pose.stomach);
        foreach marker m do
            Vector v = Vector(m - pose.origin);
            v.rotate(forward);                    // Rotate v around Z-Axis to align with forward
            pose.update(m, v);
            pose.scale(config);                      // Scale body based on human's configuration
            pose.scale(GUI.interactiveScalars());      // Adjust according to interactive scalars
            updateGUI(pose);
            RobotPacket packet = pose.strip();    // where strip() removes extraneous information
            packet.setTimestamp(timestamp / framerate);
            // Loop signal hand-gesture detected
            if handGestureDetected(LOOP) then
                packet.nextLoopState();
            end
            sendToRobot(packet);
        end
    end
end
```

Algorithm 1: Main motion-client algorithm

```
Input: framerate (Robot framerate)
// Initialization
setPose(INITIAL-POSE);                          // Set pose corresponding to human standing neutral
lastTimeStamp = 0;
while True do
    RobotPacket packet = memory.getLatestMotionCapturePacket();
    if packet.timestamp != lastTimeStamp then
        updateLoopState(packet.loopState);                 // Conveys current loop state to human
        foreach Chain c ∈ Body do
            [Θ⃗] = solveInvKin(packet.x, packet.y, packet.z);
            interpJoints(c, Θ⃗, framerate);             // Smooth interpolation of a particular chain
        end
        // Directly set AnklePitch and AnkleRoll
        interpJoints( L-ANKLE, packet.leftAnkle⃗, framerate);
        interpJoints( R-ANKLE, packet.rightAnkle⃗, framerate);
    end
    lastTimeStamp = packet.timestamp;
    sleep(1 / framerate);
end
```

Algorithm 2: Robot motion control algorithm

4.3 User Interface and Training

The *FICHMC* user interface requires very little technical knowledge to use. The training process works best with two people - one controlling the *FICHMC* graphical user interface (GUI) and one in the motion-capture suit. The first step in training is creating a configuration file for the human in the motion-capture suit. The user controlling the GUI prompts the subject to place their arms at their sides, fully extended straight forward, fully extended to the sides, and fully extended upward. This initializes our mapping scalars which correlate to the corresponding physical bounds of the robot. After initialization (or loading a previously saved MoCap config file), the human and robot are synchronized and the user controlling the GUI clicks a button to start streaming the live motion-capture data to the robot. The *FICHMC* graphical user interface can be seen in Figure 6.

The training process provides a natural interface for the human, and requires no technical knowledge at all (other than the IP address of the motion-capture server and robot). The human performs a motion, which can be seen through the real-time mimicking by the robot. The GUI includes sliders for tweaking the x, y, and z scalars of the arms and legs. For example, if the human realizes that he has to fully extend his leg (past a comfortable point) to achieve a sufficiently long stride by the robot. The x-directional scalar for the legs can be increased, making the human's forward leg motions exaggerated on the robot.

During training, it is often useful to have a "looped" motion sequence. For example, the human could take two steps and wish the robot to repeat this sequence over and over, resulting in a continuous gait. It is very important that the trainer controls when to start/stop the loop, so a simple button on the GUI wouldn't be appropriate. Extending the body-suit with an 8-marker glove on the left hand, we have implemented natural hand-gesture recognition to control the looping state of the robot. Whenever the human touches his thumb and pinky fingers together, the robot changes its looping state (and changes the color of the LEDs in its eyes to reflect to the human which state it is currently in). Optionally, the robot can also use speech synthesis to simply tell the human what state it is in through language. There are three looping states - live capture that is not being recorded, live capture that is being recorded for looping, and currently performing a loop. Every time a hand-gesture by the human is recognized,

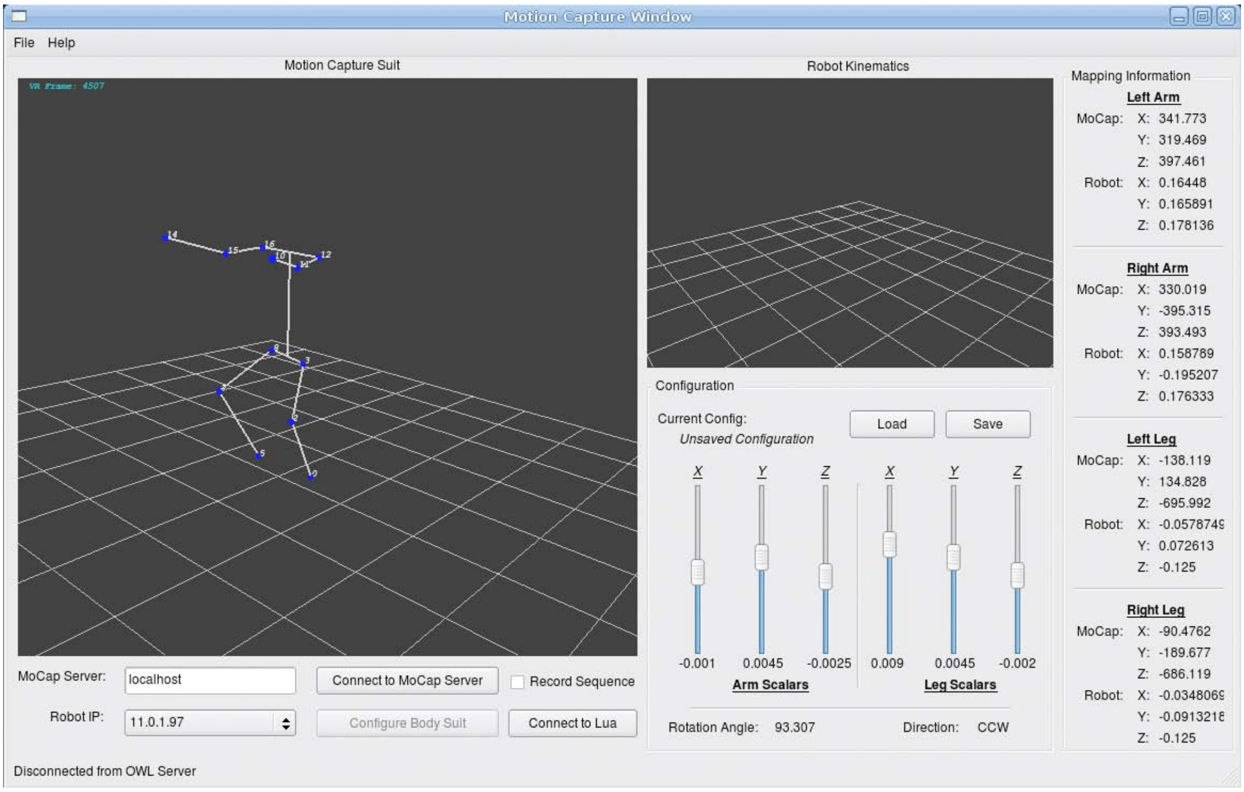

Figure 6: The *FICHMC* Graphical User Interface

FICHMC delegates the signal and changes the state of the robot accordingly. The loops are saved and can be replayed using the GUI. Additionally, the loops can be exported as a generic Lua function, that can be called from any high-level behavioral code to reproduce the looped motion.

Three images from a typical training session can be seen in Figure 7.

Figure 7: A typical training session

5 Results

Due to hardware problems in the legs of our robots, we were unable to train the Nao to walk. Although we still plan to use *FICHMC* to teach the Nao a stable and efficient gait in the future, our initial evaluation is based on upper-body tasks alone. This section primarily evaluates the ease of the training interface and assesses the ability of a human to quickly improve at a task involving fine-motor control.

Eight people were used as test subjects in this initial evaluation of *FICHMC*, and each subject completed a 45 minute interactive training session with the Nao. Our test subjects consisted of both technical and non-technical users, and each training session involved of two tasks. In the first task, a felt-tipped marker was taped to the Nao's right hand, and the robot was placed in front of a drawing board. Various shapes were put on the board (including a straight line, square, triangle, circle, "smiley-face", and the words "HI" and "HELLO"), and the user was instructed to have the robot trace over the pre-drawn shape. The second task, `Car-Park`, consisted of moving a toy car from one box to another. We discuss `Car-Park` in detail in Section 5.1.

As users practiced drawing a particular shape, their performance improved noticeably over time. Additionally, for one test, the interactive scalars were augmented at random without the subjects knowledge. There was no noticeable decrease in performance during this test, which indicates that the human was dynamically adapting its movements to those of the robot.

One user's first task was to trace the word "HI", at which they were only given one attempt. The training session proceeded normally, but their final task was to again to draw "HI". Even though they had only practiced this shape once, their results indicate a notable improvement (as can be seen in Figure 8.). This suggests that the human's learning in exploiting our interface is not isolated to a particular job, but it extends over tasks.

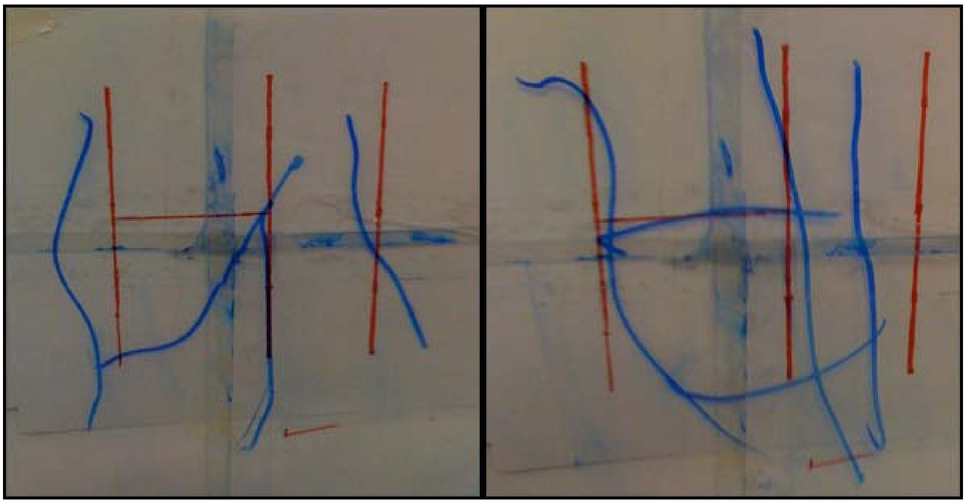

Figure 8: Tracing the word "HI" before and after a 45-minute training session

5.1 Car-Park Task

The setup of `Car-Park` can be seen in Figure 9. The robot stands in front of a surface with two distinct boxes - a source and a sink. The human stands behind the robot and attempts to puppet the robot to move a toy car from the source to the sink. The robot starts with both arms at its sides, and the task was completed when all of the car's wheels resided inside the bounds of the sink. If the car was knocked off the surface, the subject was given a 3-second penalty (in addition to the time it took to place the car back in the source).

One of the test subjects performed 60 iterations of `Car-Park`. For the first 10 episodes, the average time to completion was 28.5 seconds - for the last 10 episodes the average was 6.8 seconds. As can be seen in Figure 10, the learning curve representing elapsed time to complete `Car-Park` decreases significantly

Figure 9: The `Car-Park` Setup

over 60 iterations. This entire training session took less than 1 hour, and the subject decreased his average completion time by a factor greater than 4.

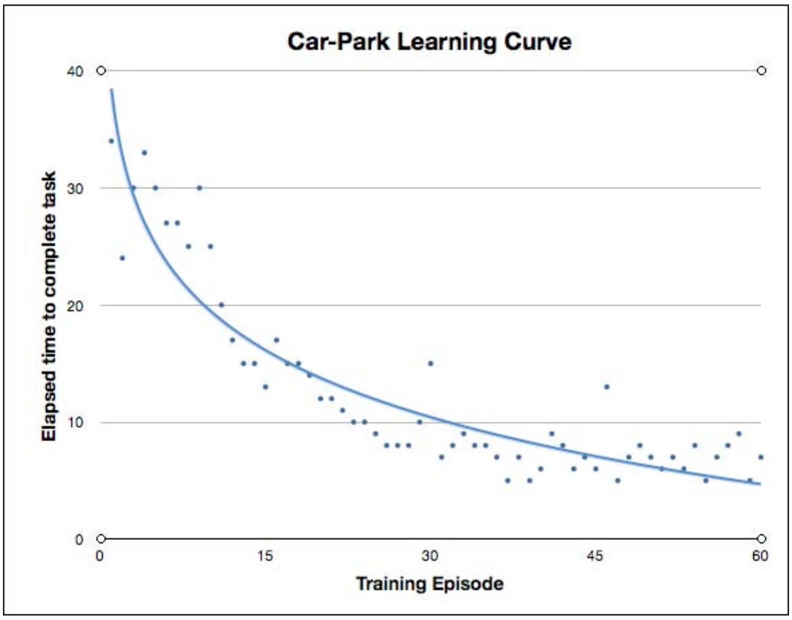

Figure 10: Learning curve obtained from 60 iterations of `Car-Park`

6 Discussion

This section contains a discussion of typical approaches towards bipedal locomotion (section 6.1). Section 6.2 presents work that is related to ours, and Section 6.3 outlines possible ways to harness the approach of *FICHMC* in future work.

6.1 Bipedal locomotion

One of our evaluation metrics will be attempting to use *FICHMC* to train an efficient gait on a humanoid robot, therefore the following sections provide a background of bipedal locomotion.

6.1.1 Passive dynamic walking

The past two decades have brought countless improvements to the field of bipedal locomotion. Prior to 1988, bipedal robotic locomotion was accomplished by applying constant power to actuators in order to completely control the robot's legs. This technique, which is largely inefficient due to constant motor output, typically results in a stiff and mechanical gait. However, in 1988 Tad McGeer revolutionized bipedal locomotion by analyzing a simple toy model of a bipedal frame. He discovered that when started on a shallow slope it could achieve a steady gait simply by exploiting gravity and the robot's momentum. He demonstrated this "passive dynamic walking" on an unpowered machine, but he argued that this breakthrough promised to lead to efficient powered walkers. McGeer's discovery started a field of robotic locomotion known as passive dynamics [22], which utilizes the momentum of swinging limbs to achieve efficient bipedal locomotion.

Since McGeer's breakthrough, research in bipedal locomotion has largely focused on exploiting passive dynamics for more power-efficient and human-like gait. In 2005, Collins and Ruina designed a passive-dynamics based bipedal walking robot that displayed a human-like morphology and gait while consuming relatively little energy [7]. This robot was specifically designed with features to exploit the nature of passive dynamics, including freely rotating hip and knee joints, direct actuation of the ankles with a spring, and wide feet shaped to aid lateral stability. These results, which built on previous work by Collins and Ruina in development of the first three-dimensional, kneed, two-legged passive dynamic walker [8], support the hypothesis that passive dynamics provides a feasible route towards animal-like locomotion.

Recent work by Owaki, Osuka, and Ishiguro [9] investigates passive bipedal running - a task where the body's intrinsic dynamics becomes increasingly dominating (due to greater momentum, there is more opportunity to exploit it through passive dynamic modeling). Their work attempts to discern an effective coupling between control and mechanics (i.e., applying power to actuators and exploiting the dynamics of the physical system). Though their research has not yet been applied to a physical robot (numerical simulation alone provides the basis of their conclusions), it was determined that the most important elastic parameters during passive dynamic running are leg spring constant and hip coil spring constant.

6.1.2 Extending the problem

Several extensions to the passive dynamic walking problem have been investigated in the past few years. For example, Asano and Luo claim that the problem has only been studied for robots without upper-bodies - their recent research investigates passive dynamic walking for a robot with a torso [13]. However, in order to preserve the natural dynamics of the biped model, they use a simple 1-link torso with a bisecting hip mechanism. Another recent paper by Ikemata, Yasuhara, Sano, and Fujimoto investigates the importance of the leg-swing mechanism in the effectiveness of an animal-like gait [14].

While we typically envision robots moving on flat surfaces, Ramamoorthy and Kuipers present a qualitative approach to bipedal walking that enables passive dynamic bipedal locomotion in realistic terrain conditions [29]. A three-phase loop controls the biped to utilize its natural dynamics. The first phase applies energy to the system, the second allows for passive leg-swing until a stopping condition is detected, and the third swaps the roles of the legs. One unique aspect of this work is that there is no limit-cycle on the number of steps the robot can take (due to the irregular nature of the terrain). This elimination of limit-cycles, which are almost always enforced in typical robotic walking scenarios, takes bipedal robotic locomotion one step closer to reaching a human level. Ramamoorthy and Kuipers extend on this work by presenting a qualitative approach to trajectory generation [30]. One interesting result stemming from this paper illustrates how to achieve task-level control over step length.

6.1.3 Alternative approaches

Passive dynamics is one blossoming approach to improving bipedal locomotion, although it is not the only method that has seen success. Pratt and Tedrake [27] compare distinct walking algorithms by factoring in the robot's center of gravity, reachable region of a single leg-swing, time of a single leg-swing, and available angular momentum. Using these parameters, along with a simple inverted pendulum model, a method for estimating the maximum stable stride-length is determined. Simulated results indicate that these estimates are beneficial using a 12 degree-of-freedom distributed-mass lower-body biped. Pratt, Chew, Torres, Dilworth, and Pratt introduce an unconventional framework called virtual model control, which is especially useful for robots without advanced sensor systems [26]. Virtual model control uses simulations of virtual components to calculate desired joint torques, creating the illusion that the simulated components are connected to the real robot. Their techniques have been successfully applied to control dynamic walking bipedal robots, assuming the robot contains foot contact switches. Finally, a truly novel approach was recently proposed by Kulic, Takano, and Nakamura that uses incremental learning of human motion pattern primitives by observation of human motion capture data [18].

6.1.4 Passive dynamics on the Nao

It is our belief that developing an efficient, human-like gait for the Nao cannot be done by simply modeling its passive dynamics for two reasons. First, the RoboCup standard platform league enforces every team to use the same exact robot - the physical specifications cannot be altered. Successful instances of passive dynamic walking normally assume a robot designed specifically for this problem. For example, Collins and Ruina [7], [8] conclude that certain physical properties are needed in order to use passive dynamics, including freely rotating hip and knee joints, direct actuation of the ankles with a spring, and wide feet shaped to aid lateral stability. Whereas the Nao indeed has wide feet, the ankles use motors for direct actuation, and the hip and knee joints are again confined by motor stiffness. In a similar vein, Owaki, Osuka, and Ishiguro [9] concluded that the most important factors in passive bipedal running are the leg spring constant and hip coil spring constant - both inapplicable to the Nao. Along the same lines, Asano and Luo [13] successfully extend the passive bipedal walking problem for robots with an upper body, but only when using a simple 1-link torso with a bisecting hip mechanism. In each of these cases, successful application of passive dynamic walking relies on certain physical properties of the robot, none of which the Nao possesses.

The second reason is that the Nao's natural dynamics are likely dominated by friction, rendering passive dynamics inapplicable. Despite the fact that the Nao has 21 degrees of freedom, each joint is directly controlled with an actuator. This joint-stiffness leads us to believe that momentum will have a minimal effect on motion.

Note that the assumptions made in the successful work of Pratt and Tedrake [27] (namely 12 degrees-of-freedom and a distributed-mass lower body) are applicable to the Nao. However, it is unlikely that the Nao's bulky legs can be modeled as an inverted pendulum. Additionally, Pratt and Tedrake's systematic approach would not allow rapid development of original motion sequences like specialized walks and kicks.

6.2 Related Work

Although motion-capture data has been harnessed to improve bipedal locomotion, to our knowledge a real-time human-robot interface using motion-capture has never been exploited in the way *FICHMC* proposes. Last year, a paper by Kulic, Takano, and Nakamura proposed a system using incremental learning of "human motion pattern primitives" by observation of motion-capture data [18]. Additionally, Nakanishi et al. have introduced a framework for learning biped locomotion using dynamical movement primitives based on non-linear oscillators, which uses motion-capture data as input [25]. While these approaches are based on a similar motivation of using human motion to train robots, *FICHMC* works in real-time, and provides a *direct* route of controlling the pose of a humanoid. Our approach may seem similar to teleoperation, but the motivation is to train robots that are able to perform motions without human instruction.

So far, we have only discussed how to represent and control bipedal locomotion in the field of robotics. However, the same basic idea is an ongoing topic in the field of neuroscience - discerning the functional basis of bipedal locomotion in humans. Last year, Azevedo, Espiau, Amblard, and Assaiante published an exciting cross-discipline paper called Bipedal locomotion: toward unified concepts in robotics and neuroscience [3]. This work claims that ongoing research in robotics can clarify our understanding of human postural control by articulating various experimental concepts and representations useful in neuroscience. On the other side, the authors investigate the inspiration behind bipedal robot design and control by human posture and gait.

Kanaoka, Shirogauchi, and Nakamura have investigated a different type of human-controlled robotics by developing a system for robotic-assisted walking using human skill and robot power [15]. Their Power-Pedal system fixes force sensors to a human's feet, and the human is able to walk around in a set of powered robotic legs. The Power-Pedal system creates a synergy of the fine motor control of humans and the theoretical brute strength of robots - the legs are able to amplify the force of a human leg by up to 40 times.

While the similarities of human movement and robotic locomotion have been investigated [3], our idea of exploiting human motor skills for efficient learning of bipedal locomotion takes a completely new approach. We will control the motion of a biped not by modeling its natural dynamics or using a virtual model of the robot, but by taking advantage of the most finely-tuned and sophisticated control mechanism known to man: himself.

6.3 Future Work

This thesis has presented a flexible framework for directly transferring human skill to a robot. We have shown that a task involving fine-motor control can be improved by over 400% in as little as 1 hour. There is still work to be done to attain our initial goal of using *FICHMC* to train the Nao to walk, but this thesis opens up some exciting possibilities for future research.

The first step in our future work will be to improve the inverse kinematics system of the Nao. In 2000, Tolani, Goswami, and Badler developed an efficient algorithm for real-time inverse kinematics specifically for anthropomorphic limbs using a combination of analytical and numerical methods [36]. The inverse kinematics implemented in $NaoQi$ uses a Jacobian approach, which can lead to more than one solution. Especially when applied to the arms, this results in the Nao often picking a solution that attempts to move the arms through it's body. While the human trainer can account for this by exaggerating its movements and exploit the mapping, improving the inverse kinematics would remove this burden.

In principle, we could adapt *FICHMC* to use a mapping from a person's arms (or even fingers) to the robot's legs. It might be that humans are more adept at rapid motor learning in the upper body than lower body - it remains an open question what type of mapping will provide the most rapid environment for motion development.

At this point, *FICHMC* restricts the learning process to the human. However, we would like to explore results in which parts of the learning process are performed by the robot. Specifically, it would be very interesting to use incremental learning in which the human performs an action, the robot attempts the action, and the human gives feedback. Note that in this scenario, both the human and the robot are "learning" - the human still learns to exploit the mapping, while the robot uses reinforcement learning with evaluative feedback to learn a model of walking. The TAMER framework introduced by Knox and Stone [16] could be adapted to work with *FICHMC*, complementing our natural human interface with a powerful system capable of evaluative reinforcement. TAMER has been shown to increase learning speed by more than an order of magnitude on the Tetris domain, and similarly to *FICHMC* requires no technical knowledge.

Last year, Kulic, Takano, and Nakamura published a system that uses incremental learning of human motion pattern primitives by observation of human motion capture data [18]. The observed data is stochastically segmented into potential motion primitives, and then a tree-based representation is built, with specialized motions at the leaves and more general motions closer to the root. *FICHMC* could combine with this method in a semi-supervised approach with an initial tutelage period to build the model of primitives, followed by training in which the robot uses live data in combination with the trained model.

If an iterative learning approach is used, we would like to make the human evaluation as natural as possible, and facilitate the ability to provide reinforcement in the motion-capture suit without having to use a keyboard. One possible way we could provide for such an easy interface is to use natural language recognition by the Nao. There have been very successful results at guiding a reinforcement learner with natural language advice [17], which could be adapted to work for the Nao. The Nao has 4 built-in omnidirectional microphones, and two stereo loudspeakers, which can easily facilitate speech recognition and synthesis. Additionally, we would likely only need the Nao to recognize two commands (such as "good" and "bad"), so training a language recognition model to understand these two words would be fairly straightforward.

Although the robot doesn't do any reasoning about the human's intentions in *FICHMC*, much work has been done on this topic. Breazeal et al. have proposed what they call *joint intention theory*, which allows robots to perform a learned task cooperatively with a human teammate [4]. While our work focuses on teaching robots how to perform tasks in isolation, it would be possible to have the robot observe the accurate representation of the trainer's actions, reason about them, and collaborate with the human.

One interesting property of *FICHMC* is that the closer the dynamics of the robot are to the dynamics of a human, the less the framework relies on the ability to exploit peculiarities in mapping the different body frames. Whereas the Nao is a very capable robot, a full-sized humanoid designed to exploit natural dynamics would likely give more positive results using *FICHMC*. As mentioned earlier, Collins and Ruina [7], [8] found that certain physical properties are needed in order to exploit passive dynamics, namely freely rotating hip and knee joints, direct actuation of the ankles with a spring, and wide feet shaped to aid lateral stability. Despite the fact that the Nao has wide feet, the ankles use motors for direct actuation, and the hip and knee joints are again confined by motor stiffness. If a full-sized humanoid (with joints not constrained by motor stiffness and direct actuation of the ankles with a spring) was combined with *FICHMC*, it is possible that the momentum of the robot could be exploited without directly modeling any of the robot's dynamics.

Although a motion-capture suit provides a very natural interface to directly control the motions of a humanoid robot, motion-capture laboratories are few and far between. In our introduction, we gave two examples of possible applications of *FICHMC*, one being remote-based surgery. A top surgeon could easily have a motion-capture lab installed near his residence, and "suit-up" every day to perform operations across the world. However, as humanoid robots become more immersed in our every-day lives, teaching a robot to perform a specific task should not require a motion-capture laboratory. Fortunately, advances in vision-based human motion capture are making great strides. A recent survey by Moeslund, Hilton, and Krüger presents some positive results. They summarize that human motion reconstruction from multiple views can currently capture gross body movement, but do not accurately reconstruct fine detail, such as hand movements or axial rotations (though motion-capture systems do this very well). While reconstruction from a single viewpoint is advancing, Moeslun, Hilton, and Krüger stress that "the use of strong a priori models enables improved monocular tracking of specific movements". They concluded that "the visual understanding of human behavior and action ... requires fundamental advances in behavior representation for dynamic scenes, viewpoint invariant relationships for movement and higher level reasoning for interpretation of actions" [23]. The current influx of research on visual-based motion reconstruction may some day allow for humanoid motion-coordination using solely mono-directional cameras.

Finally, scientists at Honda Research Institute have recently developed what they call a "Brain Machine Interface", which "uses electroencephalography (EEG) and near-infrared spectroscopy (NIRS) along with newly developed information extraction technology to enable control of a robot by human thought alone". They have yet to publish their work, but Honda has released videos of a human controlling the humanoid "Asimo" by simply thinking of which body part they would like it to move. Their system only allows for recognition of a handful of general commands (e.g., move right arm, move left arm), but has been shown to correctly identify with 90 percent accuracy. The future of human-robot interfaces may not require any movement at all...just a thought.

7 Conclusion

As stated previously, we are still working on our initial goal of training the Nao to learn a stable and efficient gait. Although our work was hindered by hardware problems, this thesis has presented strong evidence that human motor skills can be exploited to quickly train a robot. Currently, motion capture systems offer one of the easiest ways to capture the subtleties of human motion, and *FICHMC* provides a natural and real-time interface for robotic interaction, requiring very little technical knowledge.

The technical users performed slightly better during training sessions, specifically users who were familiar with the basis of inverse kinematics. The motion of these technical users was slower and more exaggerated, leading the inverse kinematics to typically choose the "correct" solution (i.e., not attempting to move the arms "through" the body). This leads us to believe that a customized inverse kinematics algorithm is needed to facilitate a more natural training environment for non-technical users.

The Nao is a capable humanoid that has served as a great test bench for *FICHMC*, but we believe that a more capable humanoid would be significantly easier to control. The Nao has no freely rotating waist and the HipYawPitch of the left and right leg are physically bound, inhibiting the ability to mimic the full range of motion provided by the human. Also, the Nao lacks direct actuation of its ankles with a spring, which has been shown to be a critical factor in mimicking humanoid locomotion [8]. A full-sized humanoid designed to exploit its natural dynamics, such as the Honda Asimo, would likely see an increased benefit by using *FICHMC*.

As mentioned earlier, *FICHMC* abstracts the task of learning away from the robot and places this burden on the human. Although our results indicate that this is a viable approach, a more dynamic set of problems could be approached if the robot and human learned in harmony. Three of the four "integral questions" for LfD proposed by Breazeal and Scassellati [5] are naturally answered by *FICHMC*, while the fourth requires the robot to reason about its actions ("How does the robot evaluate its behavior, correct errors, and recognize when it has achieved its goal?"). Our work using *FICHMC* has just begun, and we plan to address this important question in the near future.

Acknowledgements

This thesis wouldn't be possible without the vision, guidance, and perpetual help of my supervisor, Peter Stone.

Thanks to all of the members of the University of Texas Austin Villa Robotic Soccer Team, especially Michael Quinlan and Todd Hester for their unending patience and expertise. I'm extremely thankful for the opportunity to use the motion capture laboratory of Dana Ballard and John Swan-Stone - John has spent countless hours teaching me how to configure, use, and tweak the PhaseSpace motion-capture system. Also, thanks to PhaseSpace and Aldebaran, for going above and beyond typical customer service. I appreciate all the advice from members of the Learning Agents Research Group, the Reinforcement Learning reading group, and the Agents that Learn from Humans reading group.

I would like to thank my family for their continual support during my undergraduate career. And finally, I'd like to thank all of my friends here at the University of Texas. Every one of you has contributed to my knowledge, academic growth, and happiness.

References

[1] Argall, B., Chernova, S., Veloso, M., and Browning, B. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* (2008).

[2] Atkeson, C., and Schaal, S. Robot learning from demonstration. In *MACHINE LEARNING-INTERNATIONAL WORKSHOP THEN CONFERENCE-* (1997), MORGAN KAUFMANN PUBLISHERS, INC., pp. 12–20.

[3] Azevedo, C., Espiau, B., Amblard, B., and Assaiante, C. Bipedal locomotion: toward unified concepts in robotics and neuroscience. *Biological Cybernetics 96*, 2 (2007), 209–228.

[4] Breazeal, C., Brooks, A., Gray, J., Hoffman, G., Kidd, C., Lee, H., Lieberman, J., Lockerd, A., and Mulanda, D. Humanoid robots as cooperative partners for people. *Int. Journal of Humanoid Robots 1*, 2 (2004), 1–34.

[5] Breazeal, C., Scassellati, B., and LAB, M. I. O. T. C. A. I. *Challenges in building robots that imitate people.* Defense Technical Information Center, 2000.

[6] Chernova, S., and Veloso, M. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research 34* (2009), 1–25.

[7] Collins, S., and Ruina, A. A bipedal walking robot with efficient and human-like gait. In *Robotics and Automation, 2005. Proceedings of the 2005 IEEE International Conference on* (2005), pp. 1983–1988.

[8] Collins, S., Wisse, M., and Ruina, A. A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research 20*, 7 (2001), 607.

[9] Dai Owaki, K., and Ishiguro, A. On the embodiment that enables passive dynamic bipedal running. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 341–346.

[10] Dautenhahn, K., and Nehaniv, C. *Imitation in animals and artifacts*. MIT Press Cambridge, MA, USA, 2002.

[11] Denavit, J., and Hartenberg, R. A kinematic notation for lower-pair mechanisms based on matrices. *Journal of Applied Mechanics 22*, 2 (1955), 215–221.

[12] Fong, T., Nourbakhsh, I., and Dautenhahn, K. A survey of socially interactive robots. *Robotics and autonomous systems 42*, 3-4 (2003), 143–166.

[13] Fumihiko, A., and Luo, Z. Underactuated virtual passive dynamic walking with an upper body. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 2441–2446.

[14] Ikemata, Y., Yasuhara, K., Sano, A., and Fujimoto, H. A study of the leg-swing motion of passive walking. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 1588–1593.

[15] Kanaoka, K., Shirogauchi, G., and Nakamura, H. Power pedal as a man-machine synergy effectorbipedal walking with human skill and robot power. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 1779–1780.

[16] Knox, W. B., and Stone, P. TAMER: Training an Agent Manually via Evaluative Reinforcement. In *IEEE 7th International Conference on Development and Learning* (August 2008).

[17] Kuhlmann, G., Stone, P., Mooney, R., and Shavlik, J. Guiding a reinforcement learner with natural language advice: Initial results in RoboCup soccer. In *The AAAI-2004 Workshop on Supervisory Control of Learning and Adaptive Systems*. July 2004.

[18] Kulic, D., Takano, W., and Nakamura, Y. Combining automated on-line segmentation and incremental clustering for whole body motions. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 2591–2598.

[19] Lawson, A. *The neurological basis of learning, development and discovery: implications for science and mathematics instruction*. Springer, 2003.

[20] Lee, D., Kulic, D., and Nakamura, Y. Missing motion data recovery using factorial hidden markov models. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 1722–1728.

[21] Matarić, M., Eriksson, J., Feil-Seifer, D., and Winstein, C. Socially assistive robotics for post-stroke rehabilitation. *Journal of NeuroEngineering and Rehabilitation 4*, 1 (2007), 5.

[22] McGeer, T. Passive dynamic walking. *The International Journal of Robotics Research 9*, 2 (1990), 62.

[23] Moeslund, T., Hilton, A., and Krüger, V. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding 104*, 2-3 (2006), 90–126.

[24] Muellbacher, W., Ziemann, U., Boroojerdi, B., Cohen, L., and Hallett, M. Role of the human motor cortex in rapid motor learning. *Experimental Brain Research 136*, 4 (2001), 431–438.

[25] Nakanishi, J., Morimoto, J., Endo, G., Cheng, G., Schaal, S., and Kawato, M. Learning from demonstration and adaptation of biped locomotion. *Robotics and Autonomous Systems 47*, 2-3 (2004), 79–91.

[26] Pratt, J., Chew, C., Torres, A., Dilworth, P., and Pratt, G. Virtual model control: An intuitive approach for bipedal locomotion. *The International Journal of Robotics Research 20*, 2 (2001), 129.

[27] Pratt, J., and Tedrake, R. Velocity-based stability margins for fast bipedal walking. *LECTURE NOTES IN CONTROL AND INFORMATION SCIENCES 340* (2006), 299.

[28] Quinlan, M., of Newcastle (NSW). School of Electrical Engineering, U., and Science, C. *Machine Learning on AIBO Robots*. University of Newcastle, 2006.

[29] Ramamoorthy, S., and Kuipers, B. Qualitative hybrid control of dynamic bipedal walking. *Robotics: Science and Systems II* (2007).

[30] Ramamoorthy, S., and Kuipers, B. Trajectory generation for dynamic bipedal walking through qualitative model based manifold learning. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), pp. 359–366.

[31] Saggar, M., D'Silva, T., Kohl, N., and Stone, P. Autonomous learning of stable quadruped locomotion. *LECTURE NOTES IN COMPUTER SCIENCE 4434* (2007), 98.

[32] Schmidt, R., and Lee, T. *Motor Control And Learning: A Behavioral Emphasis*. Human Kinetics, 2005.

[33] Steinfeld, A., Fong, T., Kaber, D., Lewis, M., Scholtz, J., Schultz, A., and Goodrich, M. Common metrics for human-robot interaction. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction* (2006), ACM New York, NY, USA, pp. 33–40.

[34] Sutton, R., and Barto, A. *Reinforcement learning: An introduction*. MIT press, 1998.

[35] Thomaz, A., and Breazeal, C. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence 172*, 6-7 (2008), 716–737.

[36] Tolani, D., Goswami, A., and Badler, N. Real-time inverse kinematics techniques for anthropomorphic limbs. *GRAPHICAL MODELS 62*, 5 (2000), 353–388.