Accelerated and Improved
Motor Learning and Rehabilitation
using Kinesthetic Feedback

by

Jeff Lieberman

B.S. Math, B.S. Physics, MIT 2000.
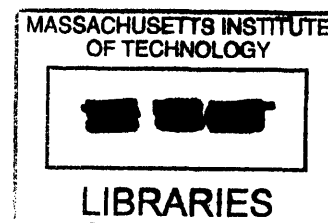M.S. Mechanical Engineering, MIT 2002

Submitted to the Department of Media Arts and Sciences
in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
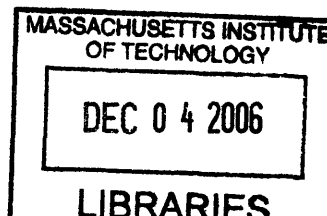
[September 2006]
August 2006

ROTCH

Author ....
Department of Media Arts and Sciences
August 11, 2006

Certified by.............................
Cynthia Breazeal
Associate Professor of Media Arts and Sciences
Director, Robotic Life Group
Thesis Supervisor

Accepted by...
Andy Lippman
Chairman, Department Committee on Graduate Students

Accelerated and Improved

Motor Learning and Rehabilitation

using Kinesthetic Feedback

by

Jeff Lieberman

B.S. Math, B.S. Physics, MIT 2000.

M.S. Mechanical Engineering, MIT 2002

Abstract

About 21 million people in the United States [roughly 8%] have a basic motor skill inability [13], many stemming not from atrophy, but an improper mapping from the brain to the motor system. Devices exist today to aid people in rebuilding their motor system mappings, but do so in bulky, and inconvenient ways, since many of the users have adequate muscle strength, but the inability to control it properly.

Hundreds of millions of people in the world participate in the arts, most of which involve motion of some sort. Typically, to become able to properly perform/paint/dance/etc, training is necessary. We learn from visual and auditory feedback, and sometimes, from the touch of a teacher. This research aims to improve the efficacy of such training with robotic touch, to enable people to become better, faster.

This research proposes an augmented sensory feedback system - a lightweight comfortable wearable device that utilizes the communication channel of direct touch on the body, to give real-time feedback to the wearer about their performance in motor skill tasks. Using vibrotactile signals to indicate joint error in a user's motion, we enable a user to wear a full-body suit that provides subtle cues for the brain, as they perform a variety of motor skill tasks. The hope is that utilizing tactile real-time feedback will act as a dance teacher or physical therapist does: by giving muscle aid through informational touch cues, not only through force or torque. This will enable people to undergo constant therapy/training, over all joints of the body simultaneously, with higher accuracy than a therapist/teacher provides.

The device will enable more rapid motor rehabilitation and postural retraining to combat repetitive strain injuries (RSIs). It will also allow allow communication between a motion expert and a student in real-time [by comparing the student's performance to an expert's], to aid in higher level motor learning skills such as sports and dance. It will function as a tool to accelerate and deepen peoples motor learning capabilities.

This thesis focuses on actuator selection and feedback mechanisms for such a suit, in a low-joint-number test, comprising elements of the upper arm. Initial tests on a 5 degree-of-freedom suit show a decrease in motion errors of roughly 21% ($p = 0.015$), with 15% lower steady-state error ($p = 0.007$) and a 7% accelerated rate of learning ($p = 0.007$).

Thesis Committee

Thesis Reader _____

Hugh Herr

Assistant Professor of Media Arts and Sciences

Director, Biomechatronics Group

Thesis Reader _____

Emilio Bizzi

Professor, Brain and Cognitive Science

Director, Bizzi Lab

Acknowledgments

The idea for this thesis came on the day the thesis proposal was due, in a dream. There was no logical genesis. I would like to thank:

My professor Cynthia Breazeal, for always trusting me to find the way, and giving the freedom to pursue the dreams in my head.

My other thesis readers, Hugh Herr and Emilio Bizzi, for providing useful feedback along the way and for making the time.

Everyone in the Robotic Life Group. They have all in some way helped along this project. Matt Berlin, Zoz Brooks, Jesse Grey, Guy Hoffman, Cory Kidd, Dan Stiehl, and Andrea Lockerd-Thomaz. Furthermore, lab alum Matt Hancher was invaluable during electronic hardware design, and taught me much along the way.

My group in Hiroshi Ishii's Tangible interfaces class, Leonardo Bonanni, Cati Vaucelle, and Orit Zuckerman, who first experimented with tactile interfaces and got my mind thinking along those lines while I was sleeping.

My family: parents Josie and Steve, and brother Chad, who are really responsible for everything I have ever done and ever will do. They are enablers of the best kind, and back it with love that makes dreams feel possible.

My friends, a limited set with limitless importance.

Finally, the artists in the world, who create experiences in life that make it one worth living.

I hear and I forget. I see and I remember. I do and I understand.

-Confucious

Contents

List of Figures

Chapter 1

Introduction

1.1 Purpose, Motivation, Applications

Roughly 8% of Americans have some motor skill inability[13]. This inability affects them on a daily basis, in everything they do. Victims of neurological trauma such as stroke form an especially interesting segment of the disabled population, because after stroke, the victims still possess adequate muscle mass, but no way to control it. The longer muscle retraining requires, the more injury can result to the body through accidents or body misuse.

On the other end of the spectrum, sports players and artists such as dancers depend on accurate motion to perform. Those that have the greatest control over their motor system perform best, and learning this control can take many years.

Typically both of these segments of the population improve their skills with a teacher - a professional who helps them improve their motor system skills. The novice/patient has three primary communication channels by which to learn new skills: auditory, usually providing high level information, or current state of performance; visual, by watching the teacher, and themselves as they attempt to perform; and tactile, through both the touch of a teacher in helping along certain motions, and in kinesthetic knowledge about one's own position/movement. A teacher cannot, however, provide real-time feedback for a novice in an efficient manner - humans have too many joints functioning in parallel, a teacher gets 'in the way' while touching a student, and a teacher uses subjective evaluation of performance, some aspects of which can be better examined by a computer vision system for accuracy.

The aim of this thesis is to minimize the time for motor skill learning or relearning in a variety of tasks, through direct muscle learning from a tactile robotic biofeedback system, as well as to remove the need for constant expert presence during the learning process. In essence, to make learning automatic and natural for anybody, to make people feel real-time feedback from an expert, from the

first time they attempt a new skill. To remove the need for an inadequate spoken language to try to teach what is otherwise so fundamental to our bodies, through the feedback of a robotic suit. We still do not know if a non-torque based tactile feedback is adequate for accelerating motor learning skills, but if it is it could represent a large shift in the way we heal patients and teach novices motor skills - compared to torque feedback, it is less bulky, less intrusive to one's behavior, can be worn unnoticeably, and requires the user to fully power the desired motion, to further train their motor system from first performance. Through human experimentation, this thesis will evaluate how well augmented tactile feedback can accelerate and improve motor learning.

This work has many potential applications. Neurological trauma and training of children with disabilities is the most immediate. Static posture analysis and retraining for those with repetitive strain injuries is a possible area for development as well. Finally, the use in the sports and arts industries is varied, from swinging a golf club, to playing guitar, to the use of a tactile signal itself as an artist medium for communication between artists. Details of these applications will be described more in depth below.

1.2 Project Scope

1.2.1 Introduction

In the piano-teacher world, it is *not* practice that makes perfect. It is *perfect* practice that makes perfect. Practicing incorrect motions during motor training can actually deteriorate users' skills, and may cause injury during the training process, as well as slow down the overall learning curve. It is very important in expert training of novices, that correct and precise motions be made properly as early as possible, before the student begins motor-learning improper motions. Feedback about performance is the single most important element of learning motor skills [7]. However, even an expert is unable to observe a student's behavior omnipotently. Furthermore, when the teacher leaves the student, performance often fades and bad habits and behaviors often supersede skills learned moments earlier. *Constant real-time monitoring and feedback of an expert*, as well as the ability to monitor *many facets of a user's performance* synchronously, is crucial to minimize the time it takes to learn a new motor-skill - quickened feedback enhances learning [35]. In fact, performance is seriously disrupted or made impossible by lags of feedback of even less than 1.0 sec [15]. We have a great ability to improve upon the typical 'this is what you did wrong' teaching method, which relies on more distant performance. Even knowing only the sign of an error in motion has been useful in accelerating motor training [20].

Teachers cannot be infallible when observing students, but they can be closer when performing the behaviors themselves. This is because they do many things of which they are no longer even conscious, through their use of *motor memory*. A teacher can quite often perform a very difficult

task, without being able to explain in words how it is done. *Spoken language is an indirect method of communicating movement - movement itself, proprioception, and touch are more direct.* Typically, the most useful task a teacher can do is correct the novice *as they perform the desired task*, by moving their muscles in the right way, with *direct tactile feedback* - this is especially true in dance and the sports arenas [e.g. golf swings]. Often times it is *not* enough for the novice to observe the expert visually and through language, as they miss subtle cues that make certain movement possible, and because direct muscle stimulation is a more direct path to motor learning than through language describing motor skills.

The goal of this system is to become a real-time, full-time, highly parallel motor skills teacher, by giving constant direct motor-system feedback to the user, as novices attempt new tasks, or patients attempt to re-attain motor skills.

1.2.2 System Implementation

The ultimate implementation concept [for eventual PhD research] is a suit worn by both an expert and an novice, as well as a software control system and motion control capture device. Its behavior is guided by the motor learning principles mentioned in Chapter 2. The system enables full body tracking of all joints, in real-time. These are recorded and stored while the expert performs the desired task [e.g. a dance maneuver], on a 3D model of the expert.

In this thesis, I will not develop the entire suit, only an upper arm subset for signal feedback testing. This upper arm will include both the elbow and wrist, so that the system cares for a hinge joint and a ball-in-socket joint, respectively. Since the body is made up of combinations of those joints, the ability to build this working subset should include all the engineering issues involved in the full body suit. In the future, for PhD work, should these initial tests show promise, the full suit will be developed and tested under a wider range of circumstances and environments. Details of this system are given in Chapter 3.

Several facets need to come together for this potential system to come to fruition. A diagram of the subsystem interaction is shown below in Figure [1-1]. Two users, one an expert and one a novice, enter the visual space covered by the Vicon [3] visual tracking system. This system covers an area of approximately 10'x20', with a 3D positional accuracy of 0.1mm for over 100 joints tracked at 120 frames/sec. Thus it is sufficient to model very accurate and rapid human motion.

Described in chapter 3 I have developed a software package to send a feedback signal to the human novice involved in motor learning. Initially, a proportional feedback signal will be given, with the gain settable by the amatuer user - initially it would be very weak, as their errors would be likely to be large, but as they gain more accuracy, the gain can be increased to show more minor errors in motion.

The novice's suit consists of joint-tracking technology, but also consists of an array of motor
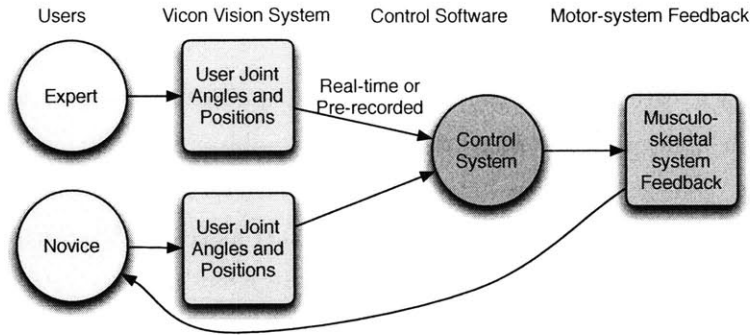
Figure 1-1: The musculo-skeletal system proposed feedback system.

output devices located at all employed joints. More detail on the specific motor technology is described below in section [3.3]. At time $t = 0$, the novice decides to attempt a performance of the new task. As they move, the tracking system observes their behavior, and compares that behavior to the expert's. Using feedback on each muscle, a signal is sent directly to the muscles, proportionally to their error from the expert's joint:

$$V_{\text{motor signal}} = K_x \left(\theta_{\text{expert}} - \theta_{\text{novice}} \right).$$

"To regulate behavior, regulate functions of error" [7]. In this manner, at any point during the performance of a maneuver, the person receives direct tactile/muscular feedback about their inaccuracies in motion. With this direct feedback over initial trials, always guiding the motion and doing so in real-time over all joints, mastery of actions is hopefully attained much more rapidly than with visual feedback alone.

Furthermore, it is possible but unknown, that this technology, once practiced by an novice, can become more immediate - they may be able to learn from muscle stimuli more quickly, and slowly accustom themselves to the point where this system becomes unconscious and forms its own feedback loop with the human's motor system. Artistically, this may extend to the point where multiple people can receive a single expert signal, and all participate in what may be a completely improvised, yet synchronous, maneuver. More importantly, this could allow a deeply shifted ability to create and retain new muscle memories.

1.2.3 Working procedure

The main engineering and research challenge is to find out what methods can influence people's motor systems, while remaining as nonintrusive as possible. Torque application systems and electrical stimulation are options, but inherently bulkier and intrusive to the user. We choose vibrotactile

stimulation, for reasons described in Section 2.4.2.

The vibrotactile actuators chosen will be tested against purely visual feedback from an expert [recorded sequences]. The most complex part of this test is deciding how the actuators will give their feedback - in what way will they signal users with a proportional error signal. Physiological research yields different possibilities, which will be tested. The test will involve a small subsection of the future bodysuit, most likely just an upper arm, in a task that requires high precision, where users will try to match a sequence of static visual depictions of arm positions [dynamic if time permits]. The tests will determine how much direct tactile feedback aids and accelerates humans' ability to learn motor positioning skills, and how much it augments pure visual feedback, to later be generalized to complex motor tasks.

1.2.4 Possible Problems

There is a large body of evidence to believe that this type of feedback system will have significant gains on muscle memory learning speeds. Notably, a teacher's direct muscle influence over a student during training is beneficial to the student, but limited in scope. However, many unknowns are involved. The method of implementation may not be ideal for the brain to internalize. Humans may lack the ability to processes so many parallel channels of communication that a system such as this requires. They may not be able to shift this learning method into an unconscious level. Because this system is so human-centric, based on as-of-yet unknown capabilities of humans, there is a lot that could not work as intended. However, all tactile input devices researched so far have shown no upper limit in informational bandwidth[26, 27, 28]. If this devices requires explicit joint attention, humans may not be able to become advanced users, internalizing so many processing channels. In this case, we will be able to incrementally increase joint information, to find the limit for humans' ability to learn such a teaching system.

1.2.5 Goals

If this system works as designed, it would represent the beginning of a marked gestalt shift in teaching methods for motor skills, across interdisciplinary fields. A measured increase in learning ability could mean permanent changes in the way we teach anyone motor skills in the future. We will be able to measure human capability to receive many parallel channels of instruction and internalize such instruction, and to determine if tactile information is enough to influence long-term behavior. Student tactile-suits could become a standard for accelerated learning in dance, sports, rehabilitation, musical instruments, sign language acquisition - in short, it could change the way we learn how to move.

Chapter 2

Background: Motor Learning, Touch, and Tactile Feedback

This chapter begins with background material that shows a motor feedback system to be both possible and promising. Feedback systems are first shown to already be a fundamental part of the way we learn motor skills. Then the nature of touch and the somatosensory system is described, indicating ways in which one can apply touch to properly feedback proprioceptive and kinesthetic information. Afterward, benefits of augmented feedback are shown. Finally previous related work is described, showing its patterns and limitations, and leaving room for further research.

2.1 Motor Learning and Feedback

Motor learning has been a subject of active research for over 50 years, and yet no deep understanding of mechanisms and methods has been found. Historically the study of motor skills learning came after World War II, when devices were developed to help Air Force pilots gain more information while flying, without requiring visual attention [such as tilt readings] [8]. As early as the late 40's [12] it was known that feedback played an important role in motor learning. Below details describe the nature of motor feedback, and its importance in learning.

2.1.1 The Importance of Feedback in Motor Learning

One point of note in agreement in the study of motor skill development is that feedback is crucial to levels of performance [6, 34, 7, 8, 9, 15, 12, 20]. Ammons [6] gives an overview of initial research done in the 1950's, noting that "The more specific the knowledge of performance the more rapid the improvement and the higher the level of performance," and that "the longer the delay in giving knowledge of performance, the less effect the given information has." Bilodeau [7] states that knowl-

19

edge of results is the "strongest, most important" variable determining performance and learning. In an extensive paper, [8], mentions that no other independent variable can affect man's ability to repeat or change his responses as feedback. Also, "performance improves immediately upon a change from delayed to immediate feedback."

Brown [12] states that feedback provides three necessary components in learning: reward, information, and motivation. Bilodeau also provides us with a highly general overview of providing feedback: "Our major conclusion on feedback is obvious: to regulate behavior, regulate functions of error."

Even the knowledge of the *sign* of the errors in a ranging task was shown to improve a gunnery trainer's performance [20]. This seemingly simple form of providing feedback based on user error determines our high-level choices in Chapter 3 regarding feedback signals.

One thing we must be careful of is the performance versus learning distinction. While wearing a feedback device, a user may perform better, but only while that feedback device is on, indicating a performance enhancement with no learning. We must be careful that this device does not increase performance with no long term gains, else it become a crutch for the user. "The transfer or shift design to separate learning from performance remains the only technique bearing on the matters of supreme importance" [8]. Section 3.4.1 describes in more detail why this system implementation shows promise to avoid this pitfall.

2.1.2 Feedback Timing

The time at which feedback is given is also extremely influential in human performance. Quickened feedback "greatly enhances" behavior and motor skill learning [35]. Conklin [15] states that performance is *seriously dirupted or made impossible* by lags of less than 1.0 seconds. We conclude that the best form of feedback is *immediate* instantaneous feedback, which allows the brain to connect synchronous actions to desired performances.

2.2 Touch Physiology

The primary touch organ, skin, is the largest organ of the human body. A thorough understanding of our sense of touch, and indeed our entire somatosensory system (comprised of the cutaneous sense, skin, the sense of movement, kinesthesia, and the organic senses), is crucial to our ability to create a viable feedback system utilizing this sense. Below we mention some of the important aspects of our somatosensory system, with regard to how we might utilize this information.

2.2.1 The Dynamics of Touch

In order to learn the best ways to touch subjects for a variety of responses, we must study the ways in which we sense different types of touch. Below we describe the primary ways in which touches are differentiated. Unless otherwise noted, this information is found in [18].

Frequency and Frequency Discrimination

The skin is very sensitive to different frequencies of applied pressure. Skin experiences its highest response to inputs of 250 Hz, and falls off at higher frequencies. Furthermore, our frequency sensitivity is sensitive to contactor size [the size of the element contacting the skin]. Smaller contactors possess a flatter curve of sensitivity across frequencies, indicating that "when small contactors are used the threshold response is independent of frequency." At higher frequencies [80-230 Hz], sensitivity increases directly with contactor size.

We respond to frequencies differently in different ranges, especially differentiating between below and above 100 Hz. "Subjects report a sensation of periodicity or 'buzzing' at low frequencies and a more diffuse, 'smooth' sensation at higher frequencies. This difference in sensory quality may be useful to the designers of tactile aids."

Furthermore, frequency content (harmonics) plays a role in tactile identification: "The vibrotactile gamut from pure sine tone to frequency-rich spectrum to noise is characterized as a continuous transition from smoothness to roughness" [33].

Our ability to discriminate between frequencies is reasonable at lower frequencies but deteriorates rapidly as frequency is increased. The skin is rather poor at frequency discrimination. Pulses produce better discrimination than do sine waves.. the difference limen for constant-frequency stimuli is better at low than at high frequencies [31].

Stimulus duration, Gap detection, Modulation, and Adaptation

The skin's ability to detect short signals depends on contactor size: "When the area of the contactor is large ($2.9cm^2$), a short signal is more difficult to detect than a long one, and ... the improvement in detectability is a very orderly function of signal duration" [38]. "Vibrotactile stimuli of duraction less than 0.1 sec are perceived as taps or jabs against the skin, providing the tactile equivalent of musical staccato." [22]

The ability to detect gaps on skin impulses exists until roughly 10 ms, but is as low as 5 ms for highly damped mechanical pulses. Bursts of sinusoids are significantly easier to detect than bursts of noise. Sinusoids are felt as 'smooth' and the gap is perceived as a small click, whereas noise feels 'rough' and the gap is perceived as a modulus of stimulus amplitude.

Similarly to the detection or application of gaps is low frequency amplitude modulation (AM). We might desire multiplication of an original signal by an LFO (low frequency oscillator) of some

type in order to indicate more information. Multiplication by a sinusoid is superior to using wide or narrow band noise [37]. Another reason to think about applying an LFO to a signal is to reduce chances of signal adaptation. Prolonged tactile stimulation can result in adaptation [22], and this is especially true when the tactile stimulus does not change over time. By applying an envelope to the signal, we can apply the 'same' signal to the skin for long periods of time, without it being forgotten by the receiver. One ubiquitous example of this phenomenon is the use of vibrating motors in cell phones and pagers. By turning these off and on at a roughly 1 Hz square wave, we never adapt to the signal and will feel it constantly.

We could also think to apply a numbered sequence of taps where the regard to number becomes important. When the number of taps is given between 2-8 Hz, the tactile capability exceeds our visual ability to count stimuli, and to regard temporal judgments of which of two signals arrived first in different locations [17]. However, for cutaneous complex displays in which many areas of the skin are stimulated in various temporal sequences, cognitive factors such as short-term memory, attention and pattern recognition become increasingly important.

When stimulus elements for sequences is increased to five or six, stimulus onset intervals needed for correct identification of the temporal sequence may be nearly 500ms. This results in a 'too slow' perception for real time speech, for example. However, to simply discriminate between two temporal sequences with no requirement to identify temporal order, increasing the number of stimulus elements has little effect on performance and discrimination thresholds are generally below 100ms. [17]

A tactile linear array was set up to test whether people can detect tactor pulse directions. "The results indicated that subjects achieved 85% accuracy in identifying the direction of tactor action." [26]

Although not as sensitive as the auditory system, the vibrotactile system can reasonably be expected to resolve temporally varying waveforms that can be utilized for processing speech information by the skin.

2.2.2   Sensory Saltation, 'The Cutaneous Rabbit'

One amazing aspect possessed by our somatosensory system is known as sensory saltation, originally found and described in [19]. It is best described with an example: We place three tactile actuators on the skin, one at the wrist, one 10 cm up the arm, and one 10 cm further. We apply 5 brief pulses to the wrist, then without any break in the regularity of the pulses, 5 more on the second actuator, and then 5 at the final actuator. Then, contrary to our default bias, "[The taps] will seem to be distributed, with more or less uniform spacing, from the region of the first contactor to that of the third" [19]. This presents a large opportunity to utilize fewer actuators to present information in a way that might usually require many more.

Several conditions are necessarily met to create this saltatory illusion. Although even 2 pulses

per location is adequate, the effect is most pronounced with 4-6 pulses per location. Any irregularity of the pulse sequence timing disturbs what has become called the 'cutaneous rabbit.' Contactors can be placed as close as 2 cm apart, and as far as 35 cm apart, while still causing the hopping effect.

Although regularity in timing is very important, the timing between taps is not highly critical. A pronounced effect occurs "over a wide range of interstimulus interval values [ISIs]" [19]. We begin to notice the effect with an ISI of 200 msec, and it settles into an evenness at 100 msec. Upon reaching 50 msec ISI, the hopping effect is optimal in regularity and vividness. "With further shortening of the ISI, the perceived [number of taps] becomes illusory. The 15 taps delivered to the three contactors mqy seem to be only six when the ISI is 20msec."

We can use this effect in multiple directions, and in fact superimpose it upon itself. "Direction of sequence is not a vital matter; hopping can go down the arm as well as up it. Indeed, it is possible to have hopping in both directions at once." Layering multiple saltatory effects on the skin at one can result "in a synergistic sum of movement on the skin" [22].

The saltation effect works equally well with electro-tactile stimulation as with vibrotactile pulses. When receiving these signals, often there is the impression that the taps "extend beyond the terminal contactor." This effect is related to 'synthetic movement' and the 'phi phenomenon' present in the visual sensory system.

2.3 Virtual Reality and Augmented Feedback

A great deal of work has been done in the last decade studying the benefits of augmented feedback, primarily given visually through a Virtual Reality (VR) environment [24, 36, 30, 11, 10]. This research can greatly inform the characterizations of feedback we apply as well as the methods of application.

2.3.1 Overview

[24] provides an in-depth review of augmented feedback in motor learning systems. A key factor of motor learning is that motor repetition is not enough to "induce cortical correlates of motor learning." The practice done by the subject must be linked to incremental success at some task or goal. Trial and error practice with feedback about performance success accomplishes this, with feedback gained through the senses.

Augmented feedback [in forms we will describe shortly] can "enhance the cortical changes associated with motor learning." Virtual reality is one methodology by which we can add augmented feedback, but none of the gains have been shown to be peculiar to VR. With augmented feedback, we receive both proprioceptive [ones sense of body position] and exteroceptive [ones sense of stimuli

outside of the body] feedback associated with the execution of a task, which "induces profound cortical and subcortical changes at the cellular and synaptic level." Visual recognition of a teacher performing a task correctly stimulates mirror neurons for learning.

Typically, the augmented feedback given is a visual display of the subject's motion, as well as a visual display of the 'correct' motion, as performed by a coach or teacher. Both motions are tracked in real-time so the user at all times can see how their motion differs from the desired motion. "VR offers the unique capability for real time feedback to the participant during practice in a very intuitive and interpretable form. Patients can see their own movement attempts in the same spatial frame of reference as that of the 'virtual teacher' (unlike practice with a real coach or therapist)" [24].

In stroke motor rehabilitation experiments, not only did motions learned in VR translate into the real world, but they also generalized to motor learning in untrained spatial locations.

2.3.2 Comparing VR with real training

Learning to perform a task consists of two primary parts:

1. Finding the set of constraints that any movement must satisfy for success

2. Selecting a subset of movements easiest to produce and control to perform reliably

These movements are known as *task related invariants* [36].

One possible way to teach task constraints is to provide reference movements that satisfy the constraints. Therefore, "one role of augmented feedback might be to emphasize the differences between the subject's movements and the reference movement" [36]. There is psychophysical evidence that humans derive specifications of movement by tracking end-effector trajectories [of the limb, usually]. By explicitly showing this trajectory through a VR display, learning may be enhanced, especially in the initial phase [10]. Furthermore, "the task can be simplified in the early stages of learning, allowing the learner to focus on key elements" [24]. Below some ways VR training may exceed natural physical training are described.

Virtual training may exceed real training

Todorov performed experiments testing table tennis strokes on untrained subjects. One group was given lessons by a human coach about how to perform the motion. A second group had their paddles' motions tracked optically, and were shown a screen with their paddle motion, superimposed on the motion of the coach who performed the task successfully. To minimize the information processing, both movements were superimposed on the same coordinate frame. "This provided an on-line error feedback to the subject during the movement in the same coordinate frame in which the subject's own paddle was being displayed" [36]. Todorov found that healthy participants who practiced

a tabletennis stroke in a virtual environment, with augmented feedback from a virtual teacher, performed *better* following training than participants who had practiced the table tennis stroke with feedback from an *expert coach*, or just practiced on their own [subjects were tested on a real world performance test]. In this case VR training taught better than a human expert!

Virtual training may show more robustness to cognitive interference

Rose et al performed a study of a complex motor skill, notably the 'steadiness tester': a metal ring that must be brought around a curved wire. They compared performance after no training, training in a VR environment, and real training. Both the real training and virtual improved significantly at the task, as would be expected. Afterwards, the task was performed alongside an interference task - the subject was required to use their other hand to tap out a tempo, in order to engage their conscious attention. The group with VR training was interestingly affected much less than the real training group [30].

Virtual training may last longer than real training

An amnesic patient was taught routes around a hospital using both real training and VR training. Upon testing two weeks later, the routes learned through VR were better remembered. Notably also, the learning was abstract, through motions of a joystick, not through walking through the halls [for the VR training] so we are necessarily learning these motor skill concepts on a more abstract level [11].

2.4 Tactile Communication and Previous Inventions

2.4.1 Related Research

Tactors [tactile actuators] were originally developed for sensory substitution, primarily for the deaf-blind community. By applying force to the skin, we can transmit coded information. The initial projects that accomplished this were such as the Teletactor (developed in 1931 by Robert Harvey Gault), an array of 32 actuators presenting sound, the Optacon (developed in the 1960s by Dr. James Bliss), a 6x24 array of actuators responding to light input [to translate written text into tactile stimulation], and the Videotact (produced in 1996 by the Unitech Company), which possess 768 electro-tactile actuators to present video [26]. These devices 'substituted' a tactile channel for the more typical auditory and visual channels one would use to process such information.

The historical development of tactile interfaces always focused on this channel substitution, relegating visual or auditory information to the somatosensory channel. This thesis focuses instead on *augmenting* the somatosensory channel with added information, utilizing its kinesthetic, proprioceptive, and labyrinthine elements to give the user a greatly added view of her behavior. It is my

belief that this historical context of sensory substitution is the only reason that the development of such an idea has taken this long. Noone has yet used a tactile interface as a sensory *augmentation* device, and never for long-term learning, only for the transfer of information. I have found no work implementing real-time non-torque-based tactile/motor feedback on users for the purpose of learning.

Below some relevant related work is described, with emphasis on how it fits into the sensory substitution/augmentation scheme.

Flight Simulation

The military recognizes the importance of muscle memory for learning applications. They provide flight simulators that allow users to learn how to use very large and expensive machinery, so that muscle memory can be learned before risking the expense of damaging such expensive machines. Companies like SimLOG [5] similarly offer purely software applications for the learning of very large construction machinery.

"In test flights with the T-34, a fixed wing Navy aircraft, pilots were able to perform simple aerobatics and basic maneuvers relying solely on haptic cues presented using a 4x5 matrix of small motors mounted in a vest" [32]. In general, haptic signals have been successfully used to convey pitch and roll information to provide pilots with aircraft attitude information [26].

Neurological Trauma Rehabilitation

Research for the aid of neurological trauma victims has led to several developments for rehabilitation, for example Myomo [4]. These devices use electromyography to get muscle signals, and apply a torque to the joint already trying to apply a torque, in order to aid weak joints. This process is known as Functional Electrical Stimulation. Their purpose is significantly different, notably applying a torque for added strength, instead of for motor skill training, especially to be learned from another. It has shown promise in neurological rehabilitation.

A patient known as HM [2] is the focus of much neurophysiological research. After brain surgery, HM was unable to form long term memories. However, he was able to form long term motor skills, known as *procedural memories*, such as bike riding or piano playing. Amazingly, he did this with *no ability to recall the training*, yet full possession of new motor skills [16], and similar patients showed the same ability to retain procedural memories, but not declarative ones. This indicates a strong likelihood that we may be able to teach people these tasks, and create long-term retention, while never making the learning methods conscious - we may be able to turn correction into an automatic muscle reflex.

Other Work

There is a great deal of other research involved in tactile aids[32, 26, 25, 18, 31] and in tactile perceptual capability [38, 37, 17]. Historically, tactile research stemmed from sensory substitution devices, beginning as a communication method for deaf-blind patients. Through decades this work has expanded to include informational displays[25, 32] of bodily status, or from external data.

Haptic displays have also been used as a balance prosthesis for vestibular disfunction, providing info about body tilt. Body sway is significantly reduced in vestibulopathic subjects using this display [25].

2.4.2 Types of Tactile Stimulation and Choice of Actuator

We have many methods by which we can apply tactile feedback to a subject. Several have been researched and are described below, including torque, vibration, and electrical stimulation.

Torque Application

Applying torque to joints is the de facto method by which to help users move their joints. However, there are several reasons to explore alternative methods. First, any motors with sufficient power to apply a noticeable torque to human joints requires a certain relatively high minimal mass, making a full body suit into more of an exoskeleton than a lightweight suit. This can impair users' motion and add unnecessary bulk to the system. A system providing *information about motion* instead of any actual muscular torque could accomplish motor learning tasks with much a much less bulky system.

Electro-tactile stimulation and EMG

The body is sensitive to electrical impulses placed upon the skin. However, for a good response, a solid electrical connection between conductor and skin is required. This typically is accomplished through a combination of shaving the receiving skin, and placing electrically conductive cream on the stimulators. For each individual, specific frequencies and voltages must be tuned to create proper tactile responses. If tuned improperly, these can cause sudden pain for the users.

Electromyography [1] (EMG) is a medical technique for measuring muscle response to nervous stimulation. It detects electrical potentials along muscle lines. Electromyography training is a kind of biofeedback in which patients learn to control muscle tension in the face, neck, and shoulders. Such training is sometimes given to migraine patients. If anything, this technique feeds back information in the form of audio, to let patients know how their muscles are responding, therefore requiring another sensory channel once again.

Electro-tactile stimulation can create unnecessary setup time, individual tuning time, discomfort, and possible pain, but is a worthwhile candidate for future research, as its main benefit is ultra-

lightweight actuators that are extremely small, which could easily be placed all over the body in a virtually unidentifiable manner.

Vibrotactile Stimulation

The use of vibrotactile stimulation, the act of vibrating some mass on the surface of the skin to stimulate a response, has shown extensive promise in many tactile aids [26]. Vibrational motors are "inexpensive, simple to control, and can produce vibrations on the skin that are readily perceptible."

[26] utilizes vibrotactile motors for a tactile vest, stating that "the frequency and amplitude variation is difficult to control independently. Feedback is typically with on/off pulsations," however utilized extremely inexpensive actuators found in beepers and cell phones. In actuators described in section 3.3, we are able to perform at a much higher bandwidth and with greater variability. Jones also attempted a shape memory alloy [SMA] version of tactile actuator, but remarked that she could only achieve a 0.2 Hz operating frequency. Given the skin's peak resonance at roughly 250 Hz this is nowhere near the bandwidth necessary for useful feedback.

For our application of motor learning, research by Matthews et al [21] indicates that application of vibration to the motor system can cause subjects to believe that their joints are in false locations due to increased muscular stress - this can be used to make the person directly perceive their errors exaggerated, without any other sort of feedback. Although not as sensitive as the auditory system, the vibrotactile can reasonably be expected to resolve temporally varying waveforms that can be utilized for processing speech information by the skin [37].

"Ideally [a wearable haptic display] should be invisible to the user until a stimulus occurs and not interfere with movements of the body" [26]. In short, vibrotactile actuators are reasonably inexpensive, able to operate quickly enough to yield useful real-time information about motor skills, can supply a large range of feedback information through variations in previously mentioned parameters, are extremely small and lightweight, and run on low power. For these reasons, the tactile feedback suit described in this thesis employs vibrotactile feedback.

Details on Actuator Selection and Usage

The science of tactile interfaces is not a large one, and therefore many open questions and applications remain. As recently as a 2004 study on torso-based tactile displays mentions "the optimal characteristics of a torso-based display in terms of the number of actuators required to present information, their spacing across the skin surface and the desired frequency and amplitude range for stimulating skin have yet to be established" [26].

Further description in section 3.3 will describe the chosen tactile actuator's capabilities and properties.

Chapter 3

System Implementation

3.1 Overview

This chapter describes the entire tactile feedback apparatus, comprised of a custom-made tactile upper body suit with optical tracking from a Vicon tracking system, 8 tactile Tactaid actuators, custom motor control hardware and firmware, and custom compare software. These eight motors regulate 5 degrees of freedom on the human right arm, shown in figure 3-5.

This suit uses motors on either side of any hinge-type joint [such as the elbow opening/closing] to give a proportional vibrotactile error signal to the joint. For example, on either side of the wrist, a vibrotactile actuator is placed. If the user moves too far in one direction, that direction will vibrate with an amplitude proportional to the angular error. However, the issue of joint rotation cannot be solved in this way. Therefore we use a sequenced vibrotactile response based on human sensory saltation, described in Section 2.2.2, in order to use the same set of motors to accomplish both hinge and rotation feedback signals. Details of each subsection are provided below.

The suit was designed and fabricated to closely model Vicon's own tracking suit, but with the allowance of inserting motors on the inside walls of the suit for vibrotactile feedback. It comprised only the right arm and shoulder subset of the larger tracking suit, to allow users an easier time getting suited and removing the suit, while experimenting.

3.2 Vicon Optical Tracking System

The Vicon optical tracking system is a commercial product designed for high-resolution high-bandwidth motion capture. It consists of roughly one dozen near-infrared sensitive cameras with matching strobes, custom hardware and custom software. A user wears a special dark suit with infrared reflectors covering it, in known locations. Each camera tracks the location of bright infrared

29
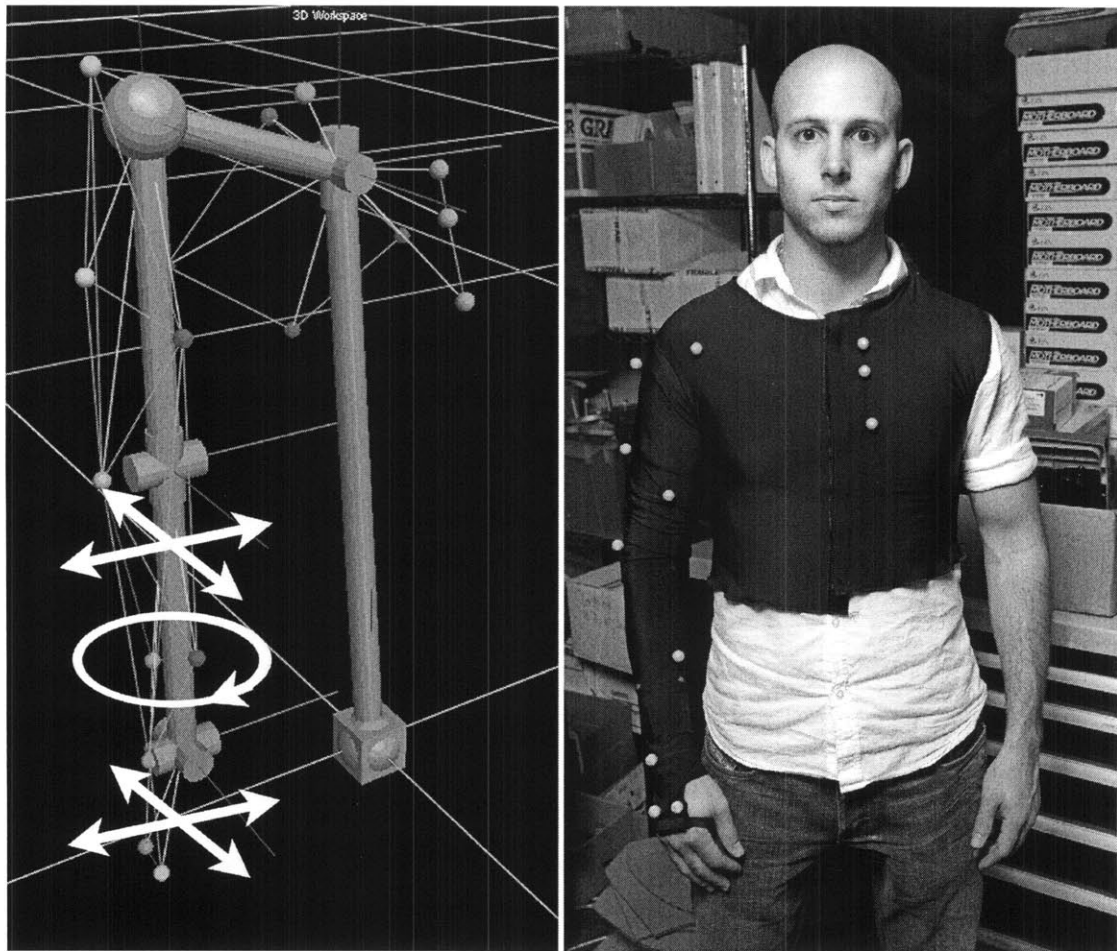
Figure 3-1: Figure indicating marker placement of the joints of the right arm, and the degrees of freedom they regulate.

reflectors, which are triangulated from multiple cameras to form points in 3D space. In order to align those with a specific body, a calibration procedure is performed, after which any movement of the suit on the body will disrupt accuracy of the results.

3.3 Tactaid Actuators

The Tactaid actuator, shown in Figure 3-2 was originally developed for speech-to-tactile translation, for the deaf community. This cantilevered resonant actuator provides extremely fast response, fast enough to translate human speech in realtime and provide it through a tactile interface, at high enough excursion [linear travel] to be felt on the skin to a reasonable degree. As noted in [14], frequency range of stimulation on the skin has typically been 230-300 Hz, regarded as optimal for contactors 7 mm in diameter. These actuators are designed to resonate on the skin at 250 Hz, the

peak frequency response of human skin. Figure 3-3 shows the frequency response of the Tactaid actuator.

Figure 3-2: The 8 tactaid actuators used in this initial experiment.
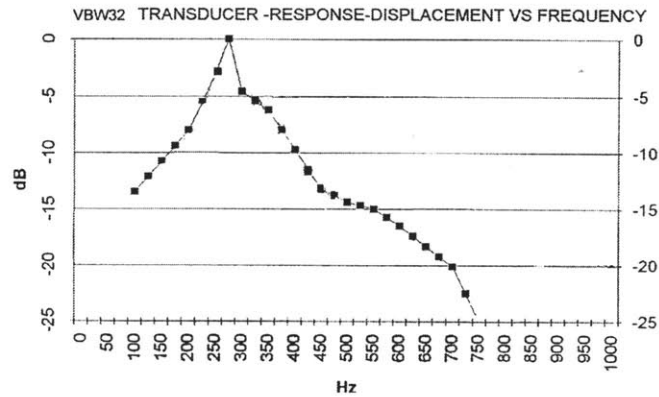
Figure 3-3: Specifications of the Tactile Actuator from Tactaid, noting frequency response.

The main advantage to using a resonant actuator of this design is that it can be turned on and off extremely quickly, thereby enabling very high bandwidth response. A typical vibrating motor in a cell phone or pager consists of a dc motor with an off-center weight attached to it. As the motor spins around, the weight is thrown back and forth very quickly generating a vibrational pulse. Two problems with this are that there is no good way to control vibrational frequency [without sensory feedback] and that the motor needs to spin up from the stopped position in order to reach the

correct frequency. Using a resonant actuator, we can completely control the frequency as it always matches our drive frequency, and as shown in Figure 3-4, the actuator rings up to full amplitude extremely quickly. We can therefore indicate signals that require high bandwidth, as is the case in many human motor skills.

VBW32 RESPONSE TESTING

Time (25 ms/div)
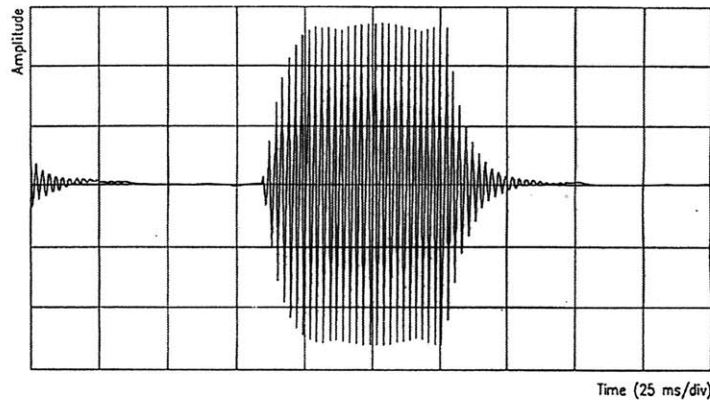
VBW32 (NOMINAL 250 Hz RESONANCE)
LOADED ON FOREARM
SQUARE WAVE MODULATION ON 250 Hz CARRIER

Figure 3-4: On-Off Response testing of the Tactaid actuator, showing ring up and ring down behavior.

These actuators are placed at the locations of both the wrist and elbow joints, in a quadrant fashion [along the major and minor axes of rotation of the joints]. In this way, we proportionally feedback specific joint angles. Slits cut into the suit allow the actuators to be slid inside, and velcro internally placed in the suit holds velcro adhered to the outside of the actuator. This insures direct actuator-skin contact, to maximize the amount of vibrations felt by the subject. It should be mentioned that no extensive testing was done to find ideal locations of these actuators, so very possibly behavior would be improved with further research into this topic. Figure 3-5 shows the locations of the 8 actuators used to regulate the 5 degrees of freedom of the right arm.

3.4 Control Software

The suit control software is written in JAVA in the IntelliJ environment on a Macintosh G5 computer. It consists of several subsystems, used to monitor teacher and student, compare motions, compute signals to be sent to the motor system, and log data for later analysis. The subsystems are described
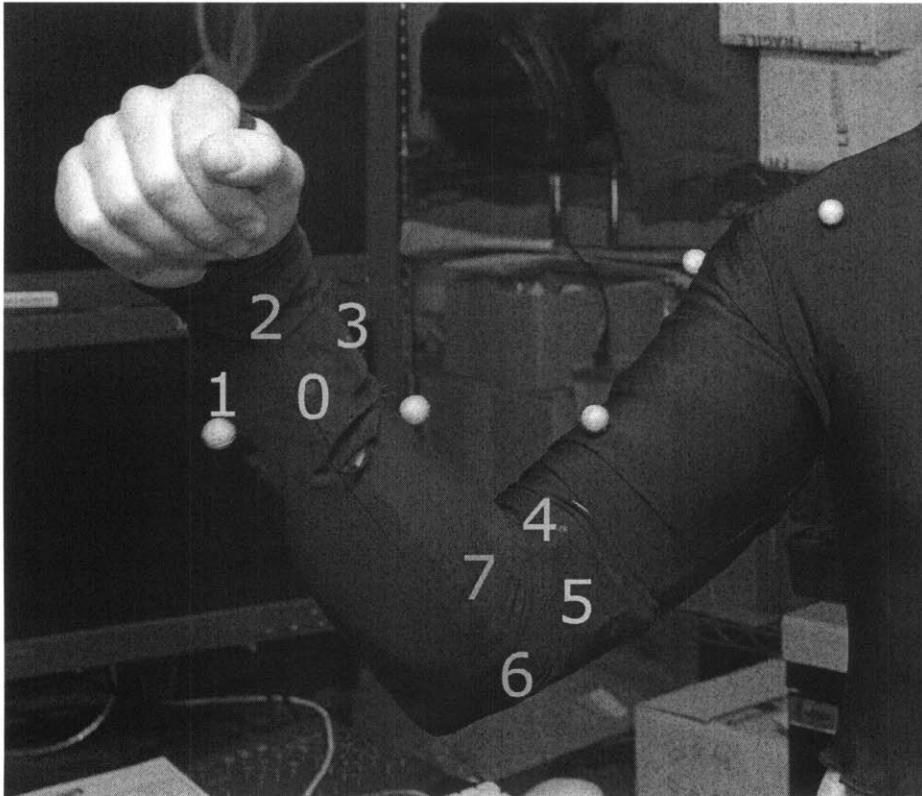
Figure 3-5: Motor placement on the tactile feedback suit. Each set of four actuators is aligned in a N-W-S-E fashion around the wrist and elbow joints, respectively. Red coloring indicates the actuator is placed out of view from this perspective, around the back.

```
Body{

    Joints[5]{
        Float measured;
        Float error;
    }

}
```

Figure 3-6: Body data structure

below.

The purpose of the software used in the human experimentation is to monitor the user's motion, while they try to perform a motion shown to them on a video screen. The original motion is captured in video form but also with the Vicon motion capture hardware/software, and is compared in real-time to the user's motions, also captured by the vicon system.

3.4.1 Code Loop

The main loop consists of the following steps. The loop occurs as fast as possible, updating video data to keep a normal video play rate, but subsampling Vicon data for higher bandwidth motion response.

Vicon Player Update [Teacher]

All video captures are stored in a data object consisting of a frame of video data, as well as a recording of all joint angles and joint positions of the right arm of the user. Joint angles are kinematicated in a manner described briefly in section 3.4.1 below. The data is stored in a Body object, which has the format

The teacher body only utilizes the measured parameter, but the student will use the comparison to generate measured errors.

Student Vicon Capture

As opposed to the teacher motion capture, the student capture must occur in real-time while giving motor feedback signals. The Vicon system continuously sends out body positions at approximately 100 Hz. The data shows optical markers 3d position, with a sub-millimeter accuracy. The Vicon system translates the optical marker positions into joint positions through a kinematic model of the users' body, but we are most interested in joint angles. Therefore we use the joint positions to determine the angles by Kinematicating them, using code developed by Zoz Brooks in the Robotic Life Group lab, similar to earlier work kinematicating motion from a telemetry suit in order to

34

operate our robot Leonardo [29].

Error Generation

We then compare the angles from the teacher with the angles from the student, and compute an error with a generalized proportional feedback system:

$$\Delta\theta_{error} = K_p(\theta_{teacher} - \theta_{student}),$$

where K_p is a constant of proportionality chosen to match the comfort level of the user. A default level of one yields very little error signal [as these angles are in radians], but since some users will by default have higher motion errors or less tolerance for high feedback signal, this is allowed to remain variable.

Note that the signal is chosen to always show error from a reference motion. Evidence from many experiments described in Section 2.3.2 shows this to be a promising approach.

It is worth noting that errors in angle are only one of many options for generating error signals. Errors in end effector/ joint position, or a hybrid between these two options, may in general lie in higher accord with the true mapping between humans when teaching motion. To be more concrete, when a very tall person and a very short person perform an 'identical' dance routine, it is unknown whether it is joint angles or positions are the most salient feature by which to measure accuracy. Most likely it is a dynamically shifting combination of the two, with possible other unknown parameters. Until this is studied, however, joint angles yield a good first approximation. Also, this study will allow more accurate measurements in the future to find those features out.

The one special case that does not represent a linear angular error is the rotational error of the wrist. This will be generated with a saltatory signal, described earlier in chapter 2. Therefore that signal is sent to the motor boards in two special channels designated as saltatory signal channels.

Motor Command Generation and Output

As described in chapter 3.5, each motor channel is given a 16 bit integer for PWM duty cycle [strength] and one bit for direction. To generate this, we first clip the signal into the range [-1,1], and then shift this float by a factor of 2^{16} and floor it to put it in the range $[-65535, 65536]$ as an integer, desired by the motor control system. We then generate the motor control packet, which consists of the format:

```
0xA5 0x28 0x00
```

$DUTY_0$ $DUTY_1$ $DUTY_2$ $DUTY_3$ $DUTY_4$ $DUTY_5$ $DUTY_6$ $DUTY_7$ $SALT_1$ $SALT_2$

```
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 CHECKSUM,
```

where the variables indicate the following parameters:

Name	Description
DUTY0	Wrist Front Motor
DUTY1	Wrist Left Motor
DUTY2	Wrist Back Motor
DUTY3	Wrist Right Motor
DUTY4	Elbow Front Motor
DUTY5	Elbow Left Motor
DUTY6	Elbow Back Motor
DUTY7	Elbow Right Motor
SALT1	Wrist Saltation Forward Signal
SALT2	Wrist Saltation Backward Signal

The $0xA5$ initiates the signal to the motor control boards, and $0x28$ signifies that we are sending desired positions/duty cycles for the motors. The checksum at the end rids us of communication errors.

Note that since all motors and saltatory signals work as opponent processes, the opponent motors should never both be on at the same time, except in the case of superposition of a linear error and a saltatory signal.

Data Logging

Every aspect of information measurable during this process is logged into a user study file, labeled by date but with no personal reference to the participant. Since the experiment consists of looping videos a set number of times, but doing so over roughly 20 different videos and still images, the master set of data is stored in a larger object format: In this manner, each video set can be repeated as often as possible, and any number of videos can be shown, while still stored in this simple data format. The master parameters store what type of feedback the user is receiving, and every frame stores all of the relevant info about joint positions and angles. All of the data is stored in a JSX object format, serialized so that we can later open this object and format it for data analysis. Also, the JSX format allows a human-readable serialized output so that if for some reason, the serialization breaks down, we can still view the file and extract the relevant data. Given the amount of data stored in this experiment every frame, the data log files are many megs and the overall data structures are several gigabytes worth of raw information.

```
MastersUser{
        boolean hasTactileFeedback;
        boolean hasVisualFeedback;
        boolean sexIsMale;
        Array videoSets[]{
                int videoNumber;
                int videoLength;
                Array videoTrial[]{
                        Array dataFrames[]{
                                int FrameNumber;
                                float liveJointAngles[5];
                                float recordedJointAngles[5];
                                float angularErrors[5];
                                Vec3 liveBodyPositions[5];
                                Vec3 recordedBodyPositions[5];
                        }
                }
        }
}
```

Figure 3-7: Data capture structure

3.5 Control Hardware

3.5.1 Motor Control Board Hardware

For this project as well as for general use in robotics projects, I have developed a motor control hardware system that allows one computer to control many independent degrees of freedom [actuators] very simply, with the motor control hardware handling all of the low level commands and feedback systems. The subsystems of this motor control hardware are described below.

Introduction

The motor control hardware system is a modular, 8-channel, dc motor control system, allowing independent feedback control of up to approximately 60 Volt motors, with peak currents of roughly 2A each. It is comprised of the following subsystems: voltage regulation (Fig. 3-9), an 8-channel analog and digital sampler, a serial computer interface, a feedback control system, and a low-level output coprocessor. These are described in more detail below. All of these systems are based on earlier work designed by Matt Hancher in the Robotic Life Group[23], remade with up-to-date parts and with a slightly different power and code architecture.
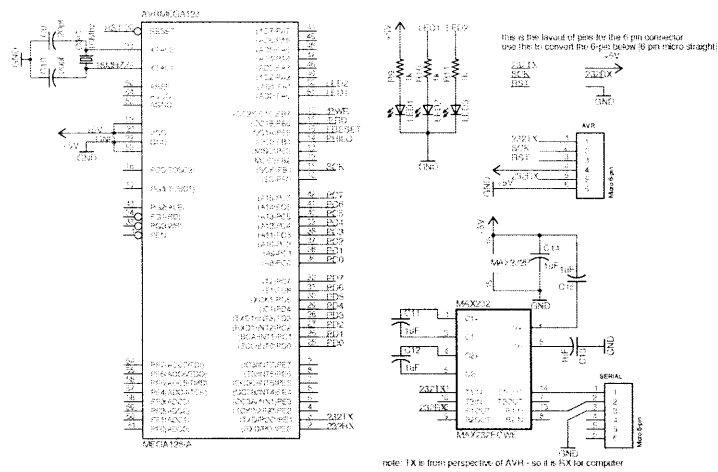
Figure 3-8: The AVR Microcontroller and communications circuitry.
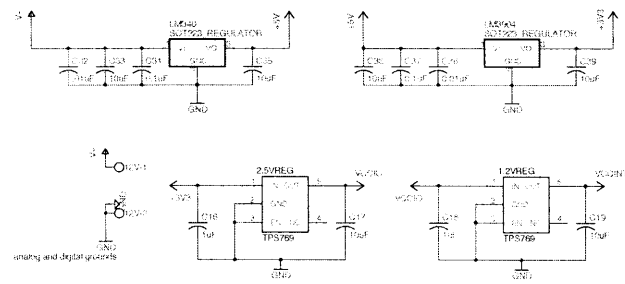
Figure 3-9: Voltage regulation and generation circuitry.

Analog and Digital Sampling

In order to get accurate positional feedback, we sample all 8 channels at 100kHz and use 16x oversampling to get rid of a factor of 4 of electrical and mechanical noise. The ADS8320 allows this high sampling rate with 16 bit accuracy. A multiplexer allows 8 channels to be sampled independently without the need for 8 expensive A/D converters. A Xilinx Spartan FPGA does all the low level co-processing - sampling the 8 channels, keeping running counts of those [and time averaging for lower noise], to be sent to the AVR whenever a request is made during the control loop.

Figure 3-10: The FPGA coprocessor and its communication channels [board-to-board connectors]

Any vibratory signal generated is the product of the high frequency (roughly 40kHz) PWM to give duty cycle, multiplied by a 250Hz square wave, to generate the vibration in the resonant response frequency of the skin. The duty cycle determined by the software is sent to the AVR which generates the 250Hz square pulsed version, which it sends to the FPGA - the FPGA then pulses the higher frequency content and sends it out via 4 channels to the H-bridge controllers. Each controller pulls up supply voltages to drive the 4 MOSFET chips, in a normal H-bridge configuration.

3.5.2 Motor Control Board Software

The software driving the AVR on the motor control board is written in C and programmed onto the AVR with an AVR-ISP. All coprocessing that is performed on the FPGA will not be described here

39

Figure 3-11: Analog inputs for position detection, force detection, and overheating detection.
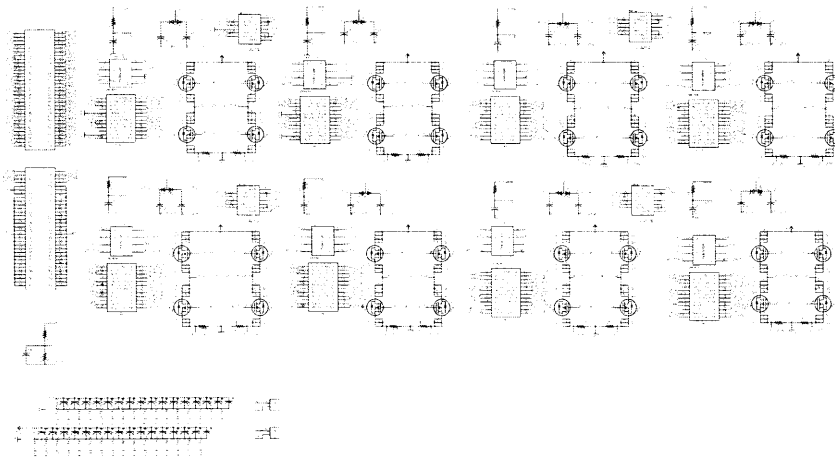
Figure 3-12: Motor output circuitry, including H-Bridge drivers and MOSFET outputs.
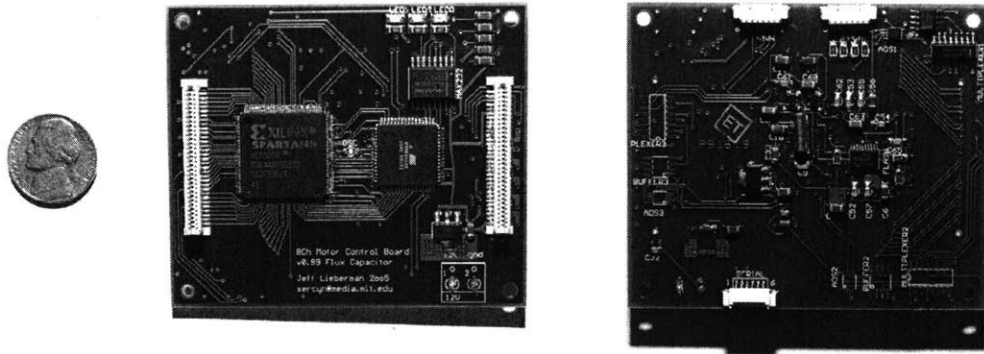
Figure 3-13: The lower control hardware board, comprised of AVR microcontroller, FPGA coprocessor, and multiplexing, a/d conversion, and serial communication hardware.
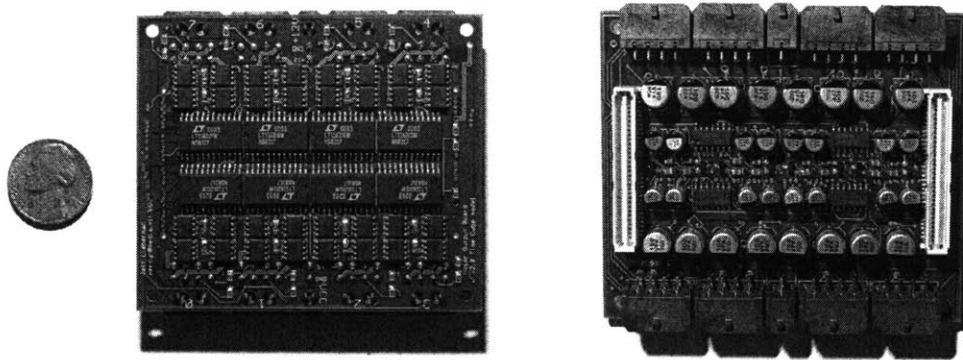
Figure 3-14: The upper control hardware board, comprised of H-bridge control modules, power i/o, and MOSFET style H-bridges, for each of 8 output channels.

in detail, as it is well covered in [23], so only small references will be given.

Upon startup, the AVR sets up a 16-bit internal timer interrupt script. The oscillator runs at 11.059Mhz, an ideal rate for serial communication [low percentage error], so by executing an interrupt every 11059 clock cycles, and using this interrupt to trigger our main loop, we are insured to have our main loop occur at a 1kHz frequency. Since we desire a perfect 250Hz output signal, as this is the resonant frequency of the skin, an even multiple of this makes the rest of our coding much easier. The AVR also sets up 57.6kbps serial communication, to receive host computer signal commands.

Several things happen during the 1kHz control loop. First, we always check to see whether we have received any new host commands from the driving computer. These commands, described in Chapter 3.4, instruct not only the duty cycle of the output motors, but also send the saltation amounts [in channels 9 and 10] to be computed by the AVR and added to the original duty cycles.

We do this because the saltation signals must move from actuator to actuator more quickly than the host computer and hardware communicate. It would therefore be impossible to control saltation accurately without a roughly 100hz communication rate, which we cannot guarantee.

Since we desire a 250 Hz square wave output, we continually downsample the 1 kHz by a factor of 2 to generate a 500 Hz signal. Every 1/500 sec, we switch direction of our output signal, which creates a 250 Hz output.

We then subsample another clock down to roughly 35 Hz for our saltation signal. This controls how often the saltatory signal jumps to the neighboring actuator. Since we use 4 actuators for our saltation, this yields a roughly 9hz saltatory rotation around the wrist. The direction of saltation is sent through channels 9 and 10, so every time this subsampled clock recycles to 0, we either increment or decrement the desired chosen wrist actuator, to carry it circularly around the wrist. There is a subsampled clock that counts half of each saltation cycle [15 of 30 cycles] and determines whether or not to turn the pulse on or off. This duty cycle influences the feel of the saltatory pulses, with a shorter on pulse yielding a more 'tap' type feeling.

Finally, the loop applies the commanded errors to each channel, and superimposes the saltatory signal on top of the original error signal. If the superimposed signal goes out of range, we clip it to its maximum value [-65535,65536]. Then this duty cycle value and direction signal are sent to the FPGA for low level PWM generation. The FPGA, as further described in [23], receives this duty cycle, and then generates a roughly 40 kHz PWM signal which is then amplified to motor level by the H-bridge circuit described above.

Chapter 4

Experimental Setup and Results

We wish to find out if the addition of tactile feedback to the normal visual feedback creates a statistically significant change in subjects' ability to learn motor skills. This change could be represented in both subjects' errors in trying to recreate motion, or in subjects' ability to improve in their performance over a long time scale. In order to do this, we test 40 subjects in a variety of motor skills tests, performed with the right arm. They are shown a sequence of images and videos showing right arm motions performed by myself, and attempt to copy those as accurately as possible in real-time, while they are monitored.

We split subjects into two groups: the first 20 receive visual feedback about the tasks they should be recreating, and the second 20 receive the tactile feedback in addition to the visual feedback, from vibrotactile actuators placed along the right arm. Subjects were roughly an even split of male and female, were between the ages of 18 and 50, with normal mobility and vision, and of height and weight to comfortably fit into the spandex feedback suit. The experiment is COUHES approved. Details of experimental protocol, calibration, testing, and results are described below.

4.1 Experimental Protocol

A user is first brought into the Robotic Life lab space, and given the COUHES agreement forms, with information about the experiment, then introduced to the general area and setup location. They are instructed that they will be wearing an optical tracking suit that will monitor their positions, and that they will be attempting to recreate motions shown to them on a video screen.

After running the subjects through a calibration routine [described below in section 4.1.1], they are sat at a desk roughly 2' wide, and 5' deep. They are sat in a stool that is height-adjusted to best seat them comfortably with their elbow laying on the table. At the opposite end to where they sit is a large 23" computer screen, that shows the videos to be recreated.

The subject is told that they will be attempting two different types of motor task. In the first,

43

they will be shown frozen images [single frames] and will try to assume the same body position [with their right arm], as accurately and as quickly as possible. After several of these images are shown, they will be shown a sequence of roughly 20 different videos. Each video will be repeated six times, and during all repetitions, the subject should attempt to recreate the motion as accurately and precisely as possible, synchronously.

They are given the tactile feedback suit, and helped with fitting. While fitting the suit, the 8 tactile actuators are placed in their testing locations on the body, on the inside of the suit. They are not built into the suit so that the suit is easier to put on and remove, and are easily installed in slits cut into the suit, once it is on. The actuators are installed in the suits of all subjects, not just the ones receiving tactile feedback, so that the physical sensation of the suit is otherwise [besides the feedback itself] the same for all users.

Once the user is fully comfortable with the suit, understands all directions, and has been shown the software, they begin. Throughout all single frames and motion videos, their motion is captured in real-time with the Vicon optical tracking system. For the tactile feedback subjects, this data is used to generate feedback signals. All users are shown the videos at all times, and all users' motion data is recorded at full frame rate, along with the associated errors in all of their joints.

At the end of the experiment, the subject's sex, type of feedback, and motion data for all recordings is stored in a file to be referenced later for data analysis. They are given a questionnaire [detailed below] and then are thanked for their time and given a Toscanini's gift certificate. Users that did not receive tactile feedback are allowed to test the system to give them an idea of what the general study is about, now that all of their data has been captured.

4.1.1   System and Subject Calibration

In order to fuse multiple 2D images from multiple video cameras into a 3D representation of the reflectors, an exact known position of each camera needs to be known. The Vicon system includes a calibration toolset that allows these camera positions to be known to sub-millimeter accuracy over a many-foot range. By waving a reflective wand around the workspace, the software will infer the positions of all cameras in use [see Figure 4-1]. After this, we can locate any other markers in the workspace.

Once the camera positions are calibrated, we can calibrate a human to the system. Everyone wearing even an identical tracking suit [mine shown in Figure 3-1] will track slightly differently, since their joint lengths and sizes will differ [for example, the elbow-wrist bone length will change]. The Vicon software includes a routine for calibrating a user to a known suit. The first half of Figure 3-1 shows the location of markers on a known suit. The user is tracked while performing a Range-of-motion (ROM) test, where they move one joint at a time in a repeated fashion, and the markers are recorded. After completed, the software processes the nearest fit of the model to the data, known
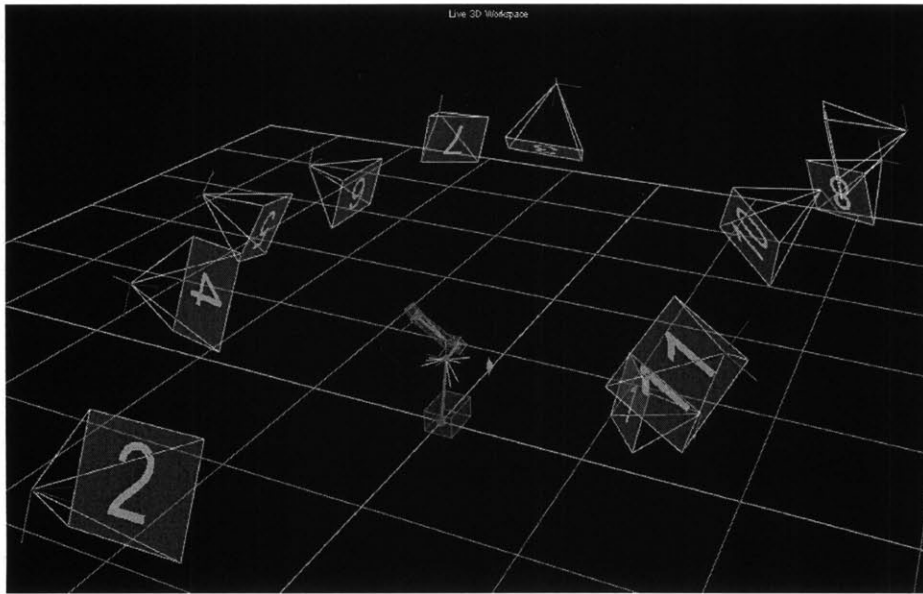
Figure 4-1: Camera Layout, determined after inital camera location calibration routine.

as a kinematic fit, and then calibrates the model by allowing adjustments to bone lengths to better fit the kinemetic data. An example of a calibrated model is shown in Figure 4-2.

After the calibration routine, they are able to see themselves being tracked in real-time on the optical tracking system, for visual confirmation that the calibration routine has allowed for sufficient accuracy in tracking. This is not a high precision test, but catches the situation in which the Vicon tracking system misinterprets points and creates large joint angular errors of roughly 90°.

After this final sanity check, the user is seated in the testing location. The setup is shown in Figure 4-3. The subject is placed seated, at a comfortable height at a desk facing a computer screen. The actuators are then connected to the driving actuator hardware. They are shown a video still of a user wearing the same suit, with the camera perspective of standing behind the user, to the left, and asked to try to put their arm in a matching position. The Vicon monitor tracking system tracks this in real-time, shown in Figure 4-4.

Once the user matches the neutral position, we zero out any linear translational errors in their angles, by clicking 'zero data' in the main user information window interface, shown in Figure 4-5. This is done to remove any first-order errors in calibration. Once zeroed, the user settings are chosen, so that they are saved with the remaining experimental data. This includes the subject's sex, and the type of feedback [visual or tactile] that the subject will be receiving. Half the users were given visual [video] feedback alone, while the other half was given tactile feedback along with the visual feedback. Once this information is entered, the experiment can begin. When the start button is pressed, the start button becomes a pause option, so that users can stop to ask questions or change their setup if uncomfortable, without capturing incorrect altered data.
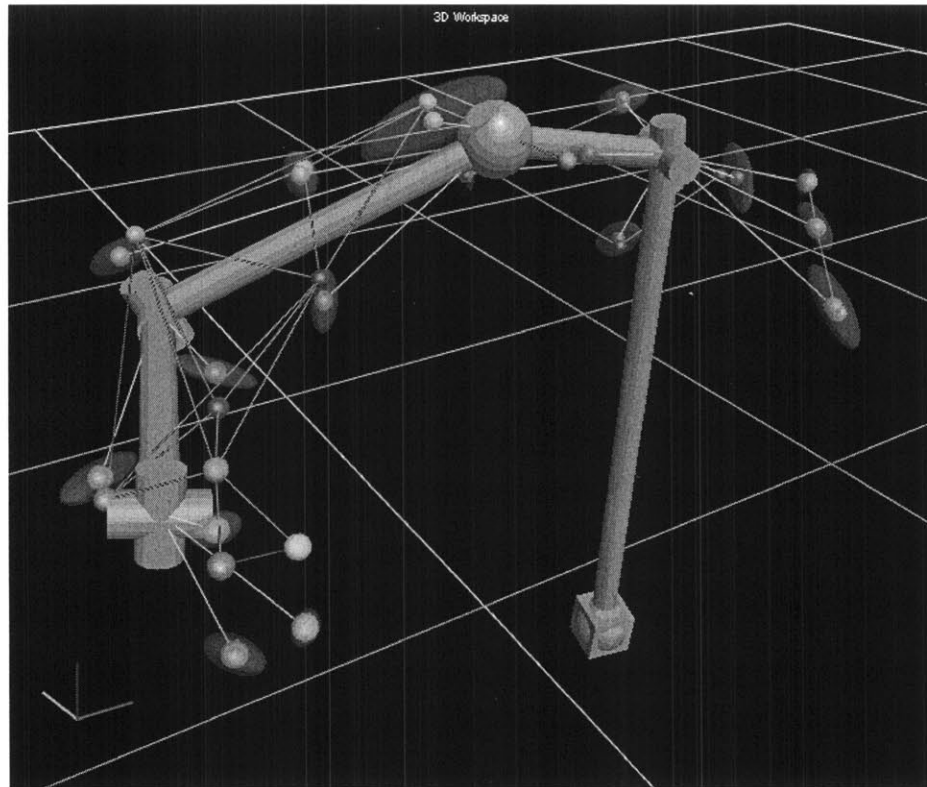
45

Figure 4-2: A view of a calibrated subject model. Translucent green ellipsoids indicate the relative uncertainty of each actuator - here, for example, the shoulder marker shows the least accuracy but the wrist degrees of freedom show very accurate response.

4.1.2 Testing Phase

The testing phase runs for roughly 18 minutes, and consists of two parts. In the first part, a series of still images is shown, similar to one in Figure 4-7, each depicting a user holding a specific position with their joint angles. This position is shown for 5-6 seconds, during which the user tries to match the position as quickly and accurately as possible. If the user is in the group receiving vibrotactile feedback, this is the first time they are presented with the vibrations, and so this is a time not only for testing, but time to allow users to accustom themselves to this new type of information. The first several images switch between a roughly neutral starting position and a position that only changes one joint angle, so that users are not suddenly overwhelmed with new vibrotactile information. As shown in Figure 4-6, all joint angles are monitored, and the 8 output signals are generated in realtime based on those angular errors.

After this roughly 4 minute initial set of tests, the subject enters a fifteen minute series of movies. Each movie is between 3 and 10 seconds long, and each is shown repeated six times. The videos show the same angle and information as the stills, but show an action progression instead of a frozen motion. While these movies play, the subjects with tactile feedback receive the feedback as they try
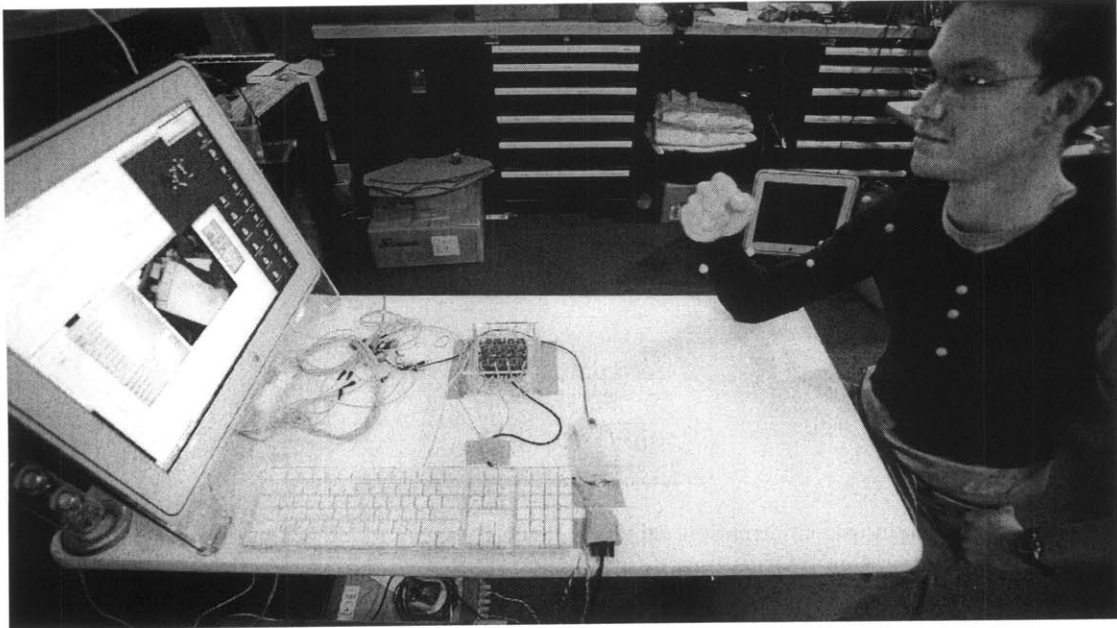
Figure 4-3: A typical user setup for the tactile feedback experiment. User is seated at a table with elbow in a fixed location, looking at the computer monitor, while being tracked by the Vicon optical tracking system.
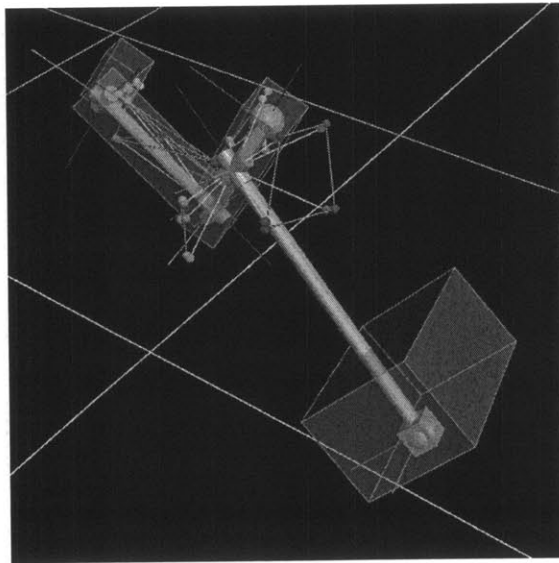
Figure 4-4: A view of the tracking system monitoring a user in progress during the experiment.

47
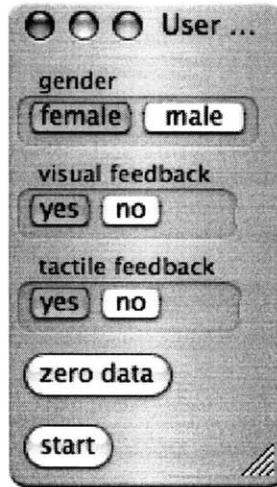
Figure 4-5: The main user information window, allowing the inputs of sex and type of feedback, as well as zeroing calibration information.
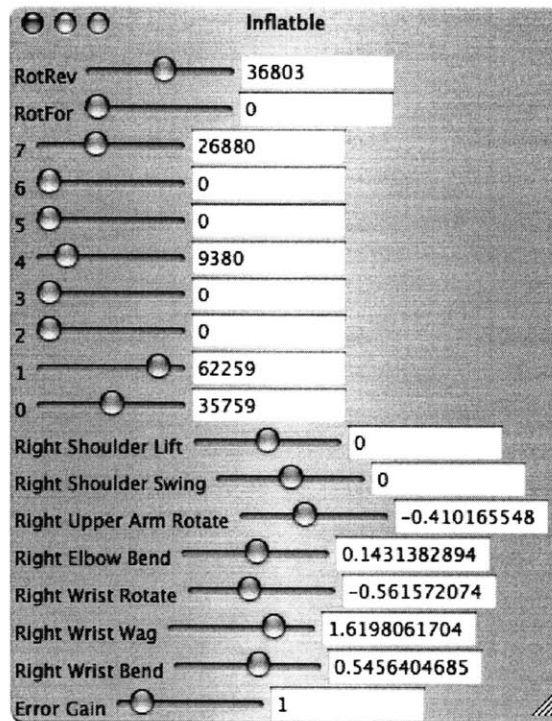
Figure 4-6: A view of the data monitoring interface, showing the 5 joint angles being tracked, and generation of the 8 motor output duty cycles for tactile user feedback.

to recreate the motions, but now this feedback is dynamic and changing with the stored videos.

The six repetitions allow the user to slowly become accustomed to the motion in the video - on the first viewing, it is unlikely that they can synchronously copy any but the simplest motions, but after six viewings their memory of previous viewings allows them to recreate motions at the same time, which much more accuracy and precision.

Through the 15 minute video series, the video actions become more and more difficult, allowing us to truly test the range of abilities over the different subjects. About 75% of the way through the progression, the videos become nearly impossible to emulate in such a short training time. Afterwards, we repeat some of the easier trials, so that we can test how much users have improved at their ability to use the system. An example of a simple motion and a complex motion are shown in Figures 4-7 and 4-8, respectively. These sets of images represent roughly 0.5 second intervals. It is easy to see in the first sequence, the motion is very slow and simple [just moving the elbow] whereas in the second, the motion is quite complex given the timing, moving all joints in many different locations.

After the physical testing is completed, the users are given a short 1 minute survey, comprised of the following statements, rated on a 1-7 scale from "Strongly Disagree" to "Strongly Agree."

- Using this device made me more fatigued than I would be in a normal motor learning situation.

- Over time I felt more capable at matching the teachers/videos motions.

- Over time I felt more comfortable with the device.

- Over time I felt more able to respond to the device.

- Using this device required conscious effort at all times.

- Tactile feedback seems a useful addition to visual feedback for motor learning.

- I was comfortable wearing the device.

- The method by which the device gave feedback aided my ability to accurately reproduce motions.

- The method by which the device gave feedback aided my ability to learn new motions more quickly.

- I received feedback about bending [ie elbow] from the device.

- I received feedback about rotation [ie wrist] from the device.

- I understand what the device is attempting to tell me.

49

1	2	3

5	5	6
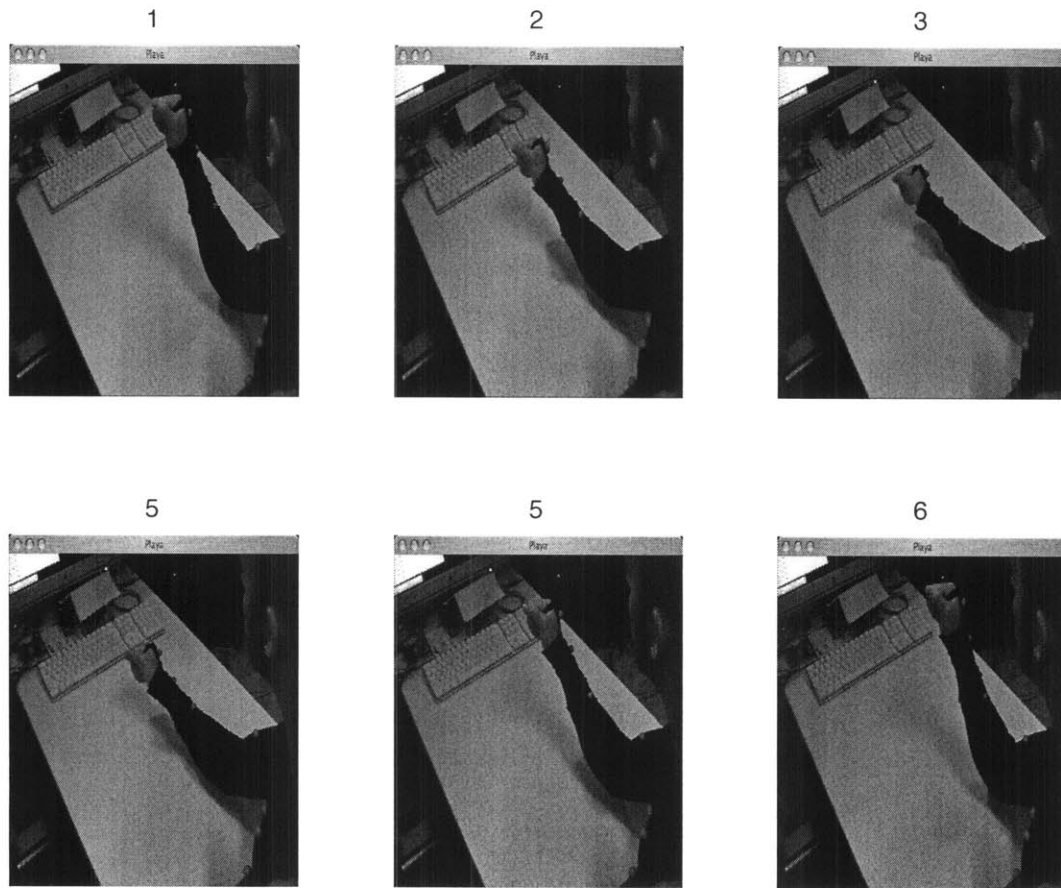
Figure 4-7: An image sequence representing 0.5sec intervals of a simple motor learning video. In this introductory video, only the elbow is moved, and very slowly, enabling subjects to get acquainted to the feedback mechanism of the system. Compare to Figure 4-8.

- I believe that if I used this device over a long period of time, I would improve at my ability to interpret information from the device.

- The method by which the device gave feedback was useful.

Users were then given a chance to enter open-ended comments about their experiences, and were then given a $5 gift certificate to Toscinini's ice cream for their participation.

4.2 Experimental Results

40 subjects were tested, 20 with only visual feedback, and 20 with visual and tactile feedback combined. After being run through the experiments, the data was analyzed in several forms; these aspects of performance are detailed below.
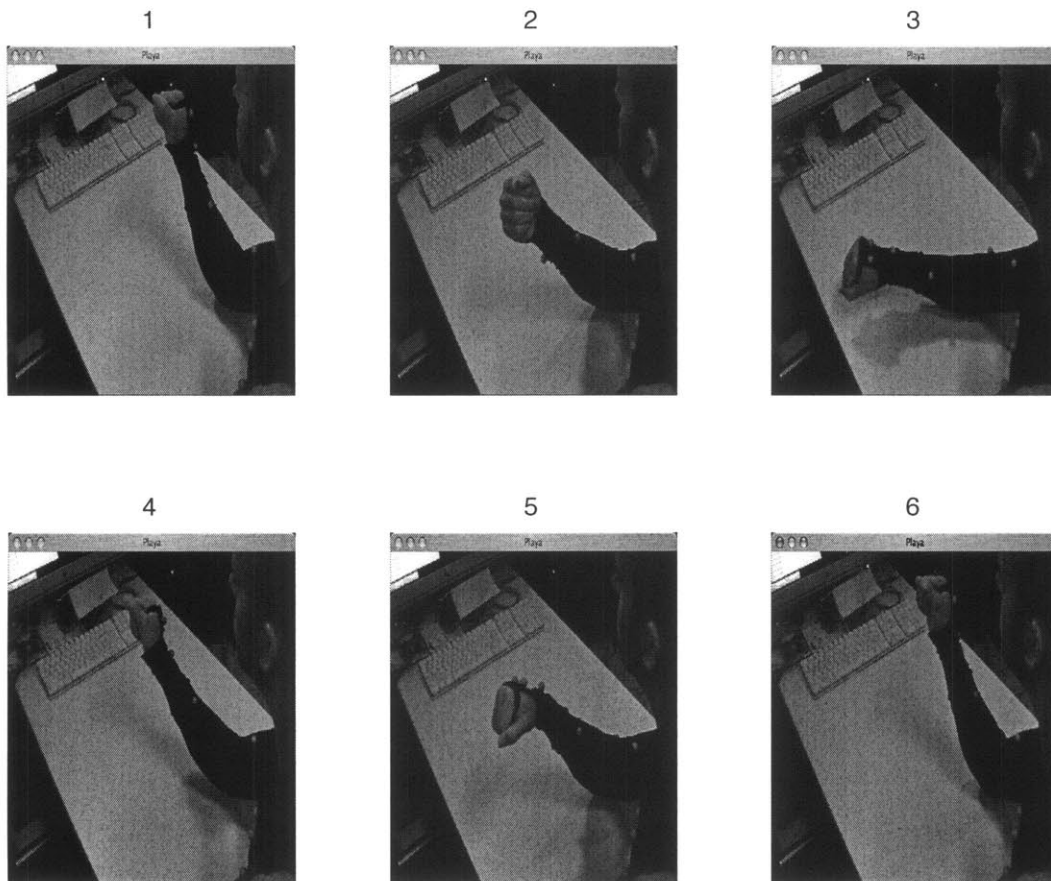
Figure 4-8: An image sequence representing 0.5sec intervals of a complex motor learning video. In this video, every joint is utilized and the motion is very dynamic, requiring subjects to use multiple repetitions to accurately mimic the video. Compare to Figure 4-7.

4.2.1 Questionnaire

There were two parts to the questionnaire, one asked of all participants, and one specific to the users wearing tactile feedback. We discuss both below. All numeric selections were based on a 1-7 scale, with 1 meaning 'strongly disagree' and 7 'strongly agree'. Below, average ratings are given in the format '(visual feedback average / visual+tactile average / overall average)'.

When asked if users were comfortable wearing the device, all users felt reasonably comfortable (5.95/5.45/5.7). Over time, they generally felt more comfortable, but the users with tactile feedback felt more comfort over time (4.74/5.55/5.15). All users felt that they grew more capable at responding to the system over time (5.43/5.75/5.62). Everyone felt they got better at matching the video's motions over time, but the tactile feedback group felt moreso (5.4/5.8/5.6). The tactile group reasonably felt that the device required using conscious effort at all times, whereas the video group felt almost neutral (4.4/5.3/4.85). No users felt that they were significantly more fatigued

than normal, but the tactile feedback group was slightly more fatigued [2.72/3.65/3.17]. Finally, all subjects felt that tactile feedback seems a useful addition to visual feedback for motor learning; but the group receiving tactile feedback felt moreso [5.63/6.1/5.89]. This was the strongest conclusion reached by all users.

The tactile group was asked some questions on their own, about the feedback system. Users felt very strongly that they would improve at their ability to use the feedback device [6.20]. Most users did not feel that the method of feedback significantly aided their abilities to learn more quickly [4.90], and they felt that the method of feedback reasonably useful [5.4]. They felt that the method of feedback aided their ability to accurately reproduce motions [5.60]. They strongly agreed that they received information about bending joints [eg. the elbow] from the device [6.35], and strongly (but less so) felt they received rotation information [eg. the forearm twist] as well [5.75]. They felt that they could reasonably understand what the device was trying to tell them [5.50].

In the free comment section, several interesting comments were made. Some users were uncomfortable with the setup: "The chair was uncomfortable", "it would be nice to have [padding] for my elbow". One user "had problems positioning [her] elbow at the right place."

Some users felt that the feedback method or amount was not ideal: "I had a bit of trouble understanding the forearm-rotate feedback signal," "the wrist rotation was a little difficult to tell what I should be correcting," "the wrist actions were more clear to understand than the upper arm area," "Sometimes there were vibrations that I didn't intuitively know how to respond to," "I frequently felt like too much feedback was coming in... times when most of the joints would be pretty close, but one or two would be providing strong feedback, and it would be difficult to tell which joints those two were." Also "the sequence is quite long, causing fatigue and distraction in the second half." One user felt that "It was hard to 'translate' the vibrations into motion commands."

Many people felt very positive about the idea in general: "very interesting research," "This is awesome!", "I like the method of using vibrations as indicators– it's nice to have stillness/ lack of vibration set in when the position and angle is correct, and I did feel that over repetitions of the same motion I felt less vibrations; I was more able to achieve the correct angle and rotation", even "I like the sounds." One user felt that it was "really fascinating; it seemed like it would be useful for the instruction of performing arts and also teaching physical movement to visually impaired people. Might also be useful in situations where there is no common verbal language."

Some had suggestions about feedback improvements: "Even when almost perfectly on, there was still some noise. I'd recommend, at least for beginning players, a larger 'dead zone' where the user is within acceptable tolerance limits. Or maybe try increasing feedback on the motor associated with the axis most wrong, and reducing / cutting feedback on the motors that are more correct."

A summary of the results of the questionnaire is given in Table 4.1.

Table 4.1: Summary of the results of the subject questionnaire; the first section was given to all participants, and the second only to those with tactile feedback. All questions were answered on a 1-7 scale, where 1='Strongly Disagree' and 7='Strongly Agree.' Significant results are highlighted in boldface.

Question:	Visual Average	Tactile+Visual Average	Overall Average
I was comfortable wearing the device	5.95	5.45	5.7
Over time I felt more comfortable with the device	4.74	**5.55**	5.15
Over time I felt more able to respond to the device	5.43	5.75	5.62
Over time I felt more capable at matching the teachers/videos motions	5.40	5.80	5.60
Using this device required conscious effort at all times	4.40	**5.30**	4.85
Using this device made me more fatigued than I would be in a normal motor learning situation	2.72	**3.65**	3.17
Tactile feedback seems a useful addition to visual feedback for motor learning	5.63	6.10	**5.89**
I believe that if I used this device over a long period of time, I would improve at my ability to interpret information from the device		**6.20**	
The method by which the device gave feedback aided my ability to accurately reproduce motions		5.60	
The method by which the device gave feedback aided my ability to learn new motions more quickly		4.90	
I received feedback about bending [ie elbow] from the device		**6.35**	
I received feedback about rotation [ie wrist] from the device		5.75	
I understand what the device is attempting to tell me		5.50	
The method by which the device gave feedback was useful		5.40	

53

4.2.2 Image Tracking

The first measure of performance was the frozen snapshots, where only a solid frame of video was shown to participants, as they tried to track the information. The average performance in tracking is shown in Figure 4-9. This is an average of every user, over every snapshot shown to them, thus it is the average performance in one showing of a frame and tracking to it.
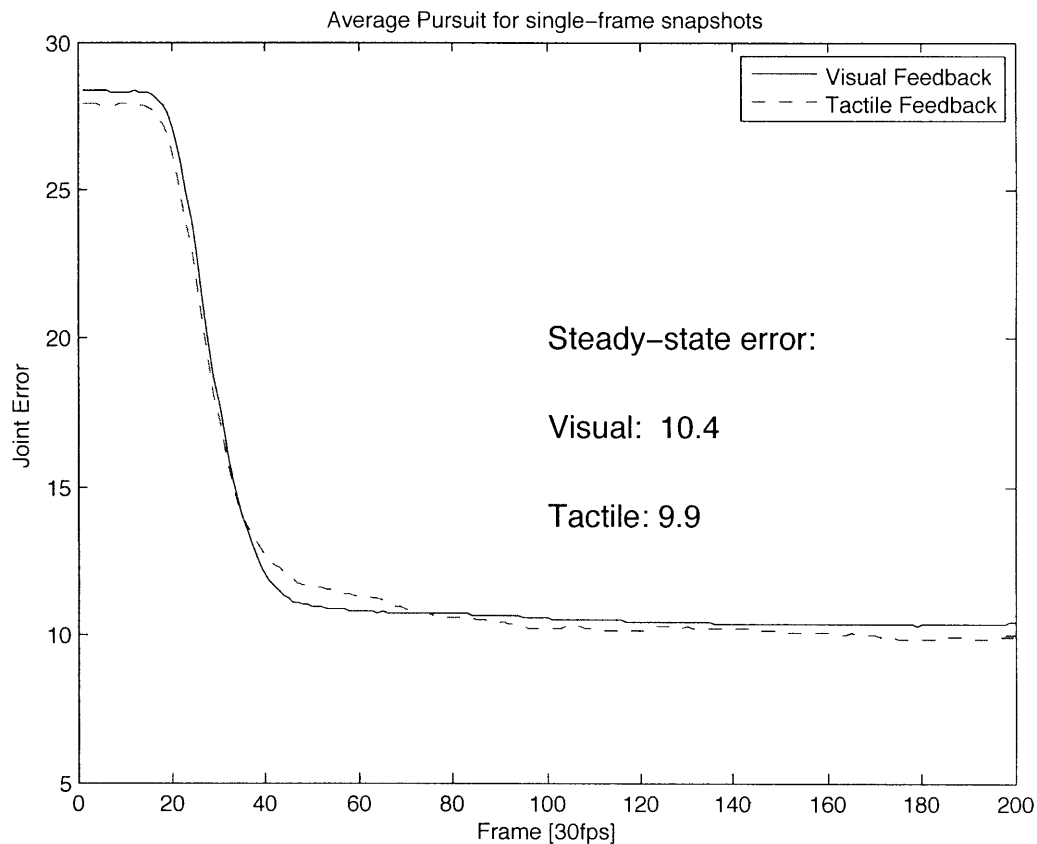
Figure 4-9: User tracking of single-frame images. Note that both types of user possess roughly the same temporal ability to track, but tactile feedback improves steady state error by roughly 5%.

It should be noted that many users receiving tactile feedback used much of this time to experiment with the user space; so, this gives a very conservative view of users' ability to track with tactile feedback, as junk data is averaged into the normal tracking data with no easy separability. Even so, over the 20 snapshot trials given to the 40 users, the tactile feedback did result in slightly better steady-state tracking performance of roughly 5%. This is not statistically significant (first frame significance $p = 0.20$, last frame $p = 0.08$), but once again includes much data where tactile feedback users were experimenting. So, this gives a worst case scenario.

4.2.3 Video Tracking

The bulk of the experiment involved subjects tracking moving videos. In order to analyze errors, a choice needs to be made between end-effector error ($< \vec{x}_{subject}(t) - \vec{x}_{goal}(t) >$), and an angular error, which we define as

$$\epsilon = \sqrt{\theta^2_{\text{wrist rotation}} + \theta^2_{\text{wrist in/out}} + \theta^2_{\text{wrist left/right}} + \theta^2_{\text{elbow left/right}} + \theta^2_{\text{elbow open/close}}},$$

where θ_x means the angular error of joint x. During the experiment it was made clear that subjects should try to mimic joint angles, not necessarily end effector position. Therefore we use the joint angle error calculation as our metric for tracking error.

It should be noted here, that at several points during several subjects, the Vicon system suddenly stopped properly tracking one or more of the joints [usually in the wrist]. Whenever this happened, the system would register a very large error that would need to be reset by pausing the program, removing the subject from the testing area, and bringing them back into the area, allowing re-registration of the subject by the Vicon system. This definitely effected the outcome data, until the problem was recognized. However, in analyzing the subject data after all subjects were tested, the data appeared very consistent over the entire 45 minute length of the experiment. Thus, it is believed that individual problems were more than covered by the number of subjects tested, and no special routines were added to remove the spurious data.

The first element analyzed was purely a measure of angular error over all time for all participants, split into two groups by type of feedback given. This is shown in Figure 4-10. Red indicates the users given tactile feedback on top of the visual feedback given to all users. In general, the tactile feedback given lowers angular error over the entire experiment.

In order to get a more quantitative view of this improvement, we analyze this data further. Individual segments of the experiment have ranging difficulty; in order to get rid of this from our view, we divide the errors in real-time of the tactile feedback group by that of the visual feedback group. This gives us a *relative error* measure for the entire experiment. Figure 4-11 shows this relative error metric.

Two notable facts emerge: first, almost the entire performance is enhanced by the use of tactile feedback. Only for very brief periods of the trials, likely the first viewing of a new motion, does any visual performance exceed that of those with tactile feedback. More importantly, over the entire set, the addition of tactile feedback greatly reduced subject error. Overall, subjects utilizing this new form of feedback on top of visual feedback are 21% better at matching novel motions presented to them. This result is extremely statistically significant ($p = 0.015$).

It is notable in Figure 4-11 that there is some regular oscillation in performance, and this is from the fact that each new movie was shown multiple times. Therefore, a further analysis is used to find
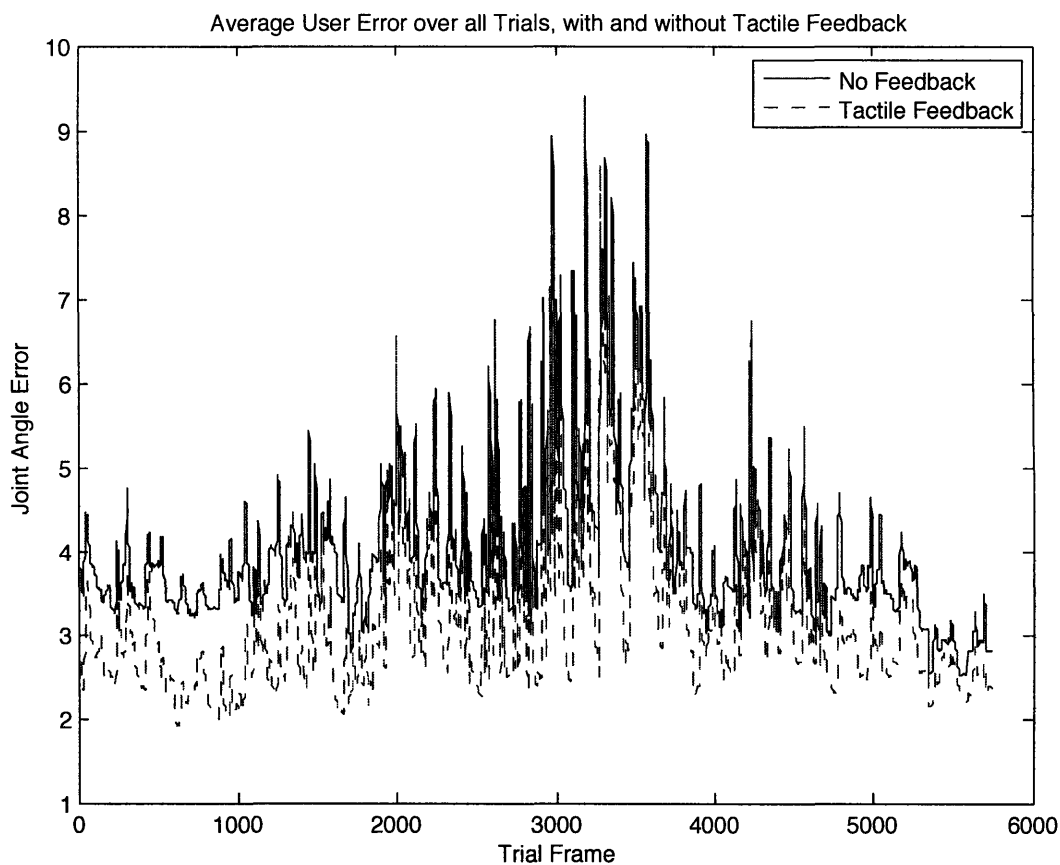
Figure 4-10: A timeline of average angular error for all subjects over the entire length of the video segment of the feedback experiment.

user's improvements over multiple playings.

4.2.4 Improvement over Trials

Each video was shown 6 times to each user, to allow users the ability to try to improve over time, and to test how users react to such feedback. Since the long-term purpose of this device is for teaching/learning novel motions, improvements over time are the most critical aspect. The above analysis shows the first part of that, the initial error. The smaller an initial error, the less a subject needs to improve to reach some criterion of success, whatever that be. However, from any starting point, one must improve at least as quickly, if not more quickly, for this device to show utility in a learning context.

Therefore, we analyze subjects' performance over the 6 trials of each video set. We study their integrated error over a trial,

$$\Delta_n = \int_{t_n}^{t_{n+1}} \epsilon(t)dt,$$

Figure 4-11: Relative error between the tactile feedback group and the visual feedback group. Numbers less than 1 indicate tactile feedback performance exceeding visual feedback. On average, the addition of tactile feedback outperformed visual by roughly 21%.

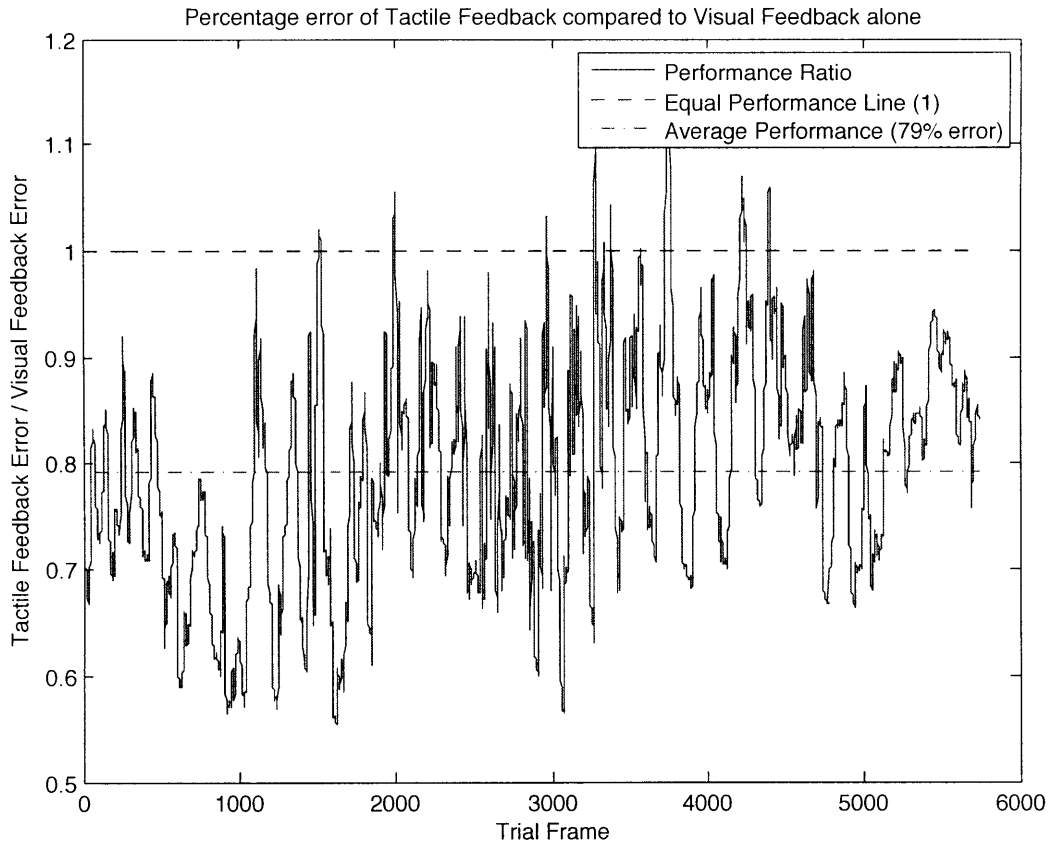where $[t_n, t_{n+1}]$ marks the time interval for a the n^{th} trial, and, as earlier, $\epsilon(t)$ represents the angular error between subject and desired position. Figure 4-12 shows the result of this analysis.

The most interesting thing about this initial look at the trials is that, at no point does the addition of tactile feedback debilitate subject performance; at all times, the addition lowers subject error. This is independent of task difficulty; at all times, performance is enhanced. To get a better idea of this, we can once again look at a fractional performance metric, which will get rid of individual task difficulties, and highlight relative performance. Figure 4-13 shows this relative performance.

At all times of this section of this experiment, almost 15 minutes, the addition of tactile feedback aids the performance of the subjects. How much and in what ways? We are interested in looking at not only subject accuracy, but also the speed at which they settle on their behavior, and the ability to improve over time with multiple trials.

The easiest way to find a quantitative answer to this is to combine all separate movie data into one. Each movie was shown six times; here, we average all of the movie's n^{th} trial together, to
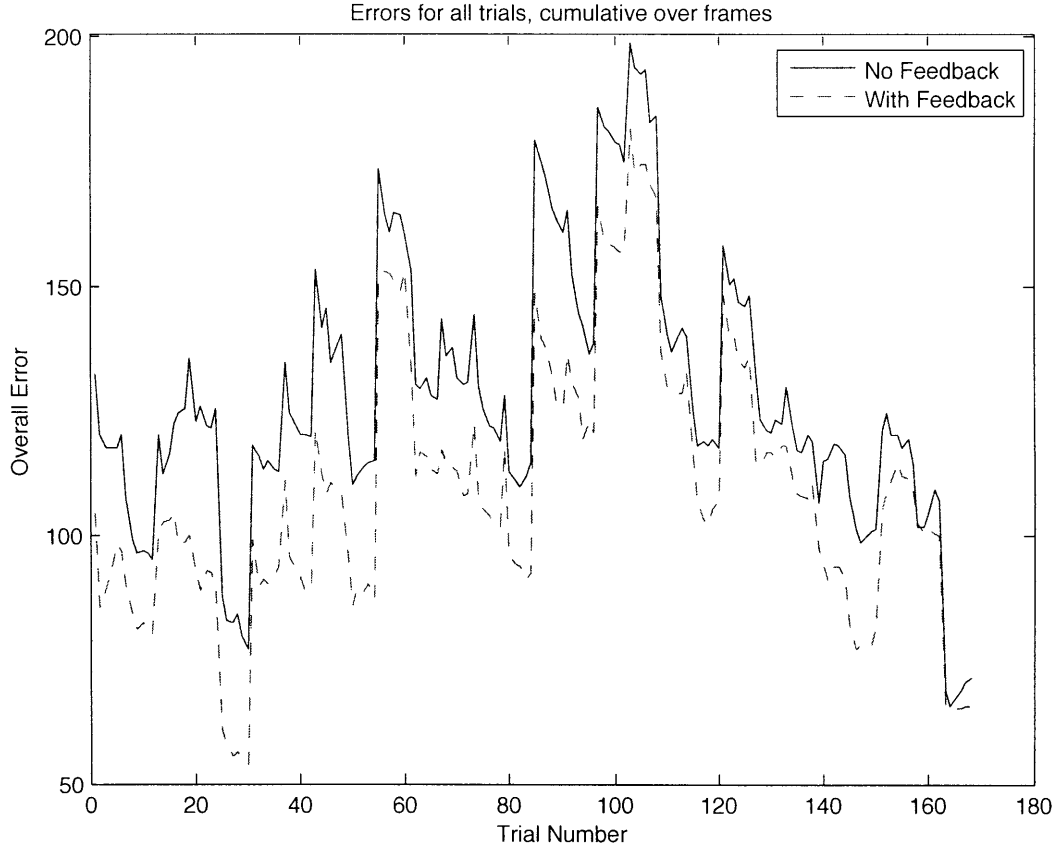
Figure 4-12: A comparison of abilities during all trials of each movie progression. Each point represents one full trial, therefore every six points represents one movie set. Note, through the entire progression, tactile feedback outperforms visual feedback.

get one overall n^{th} trial performance. Then we average this metric over all users. Overall this is represented by

$$\Delta T_i = \frac{1}{\# \text{ subjects} * \# \text{ movies}} \sum_{u=1}^{\# \text{ subjects}} \sum_{m=1}^{\# \text{ movies}} \Delta_i,$$

where u represents the different users, and m represents the different movies, with Δ_i still representing the integral in Equation 4.2.4. The result of this is shown in Figure 4-14.

For ease in analyzing this sort of data, we assume that user abilities will settle over time in a fading exponential model [which seems to fit our data points quite accurately]. This model is represented in the following functional form:

$$y = a + be^{-cx}$$

where x represents the trial number, y represents the trial error, and a, b, and c are free parameters.
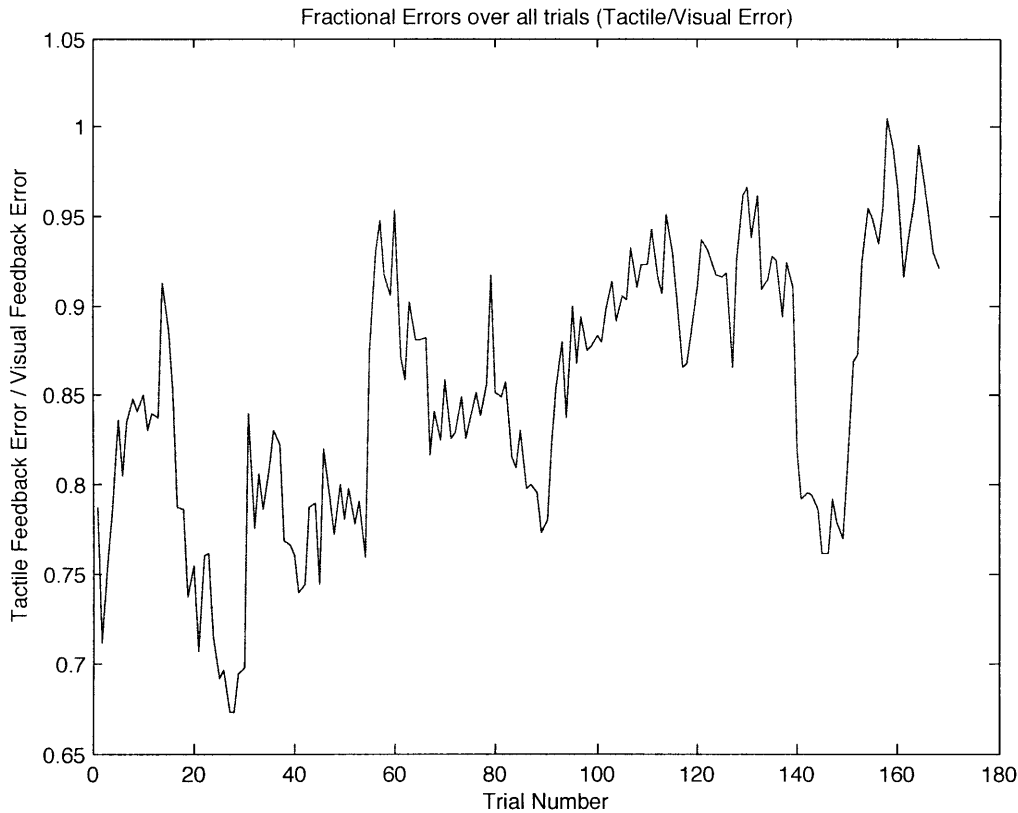
Figure 4-13: Analysis of trial error, showing relative performance of tactile + visual feedback, compared to only visual feedback. Note that the plot is always below unity, indicating a performance enhancement at all times of the experiment.

Both curves were fit using a linear least squares fitting form.

This analysis shows several interesting features. First, even on first viewing, the average subject error in recreating novel motion is reduced by roughly 11%. The parameter $a + be^{-c}$ records this initial error.

Secondly, the parameter c represents the time constant of this learning: the higher c represents the subject learning faster, settling more quickly on their final behavior. Here we see a time constant of 1.52 become 1.63 with tactile feedback, an improvement of 7%.

Finally, the parameter a alone represents the steady-state error after infinite trials. Here we see the steady-state error begin at 125.2 and lower to 106.9 after the addition of tactile feedback. This represents a *reduction in steady-state error* of 15%.

In order to gauge the statistical significance of these findings, we find the statistical significance of each trial independently. For each trial, all frame errors are averaged for each user, and those averages determine the distributions of both populations. In each trial, the results were consistently extremely significant: in order of trial, the significance values are $p = 0.007, 0.007, 0.009, 0.007, 0.004,$

59

Subject Error over repeated Trials

$$y = 125.2 + 45.9e^{-1.52x}$$

$$y = 106.9 + 46.1e^{-1.63x}$$
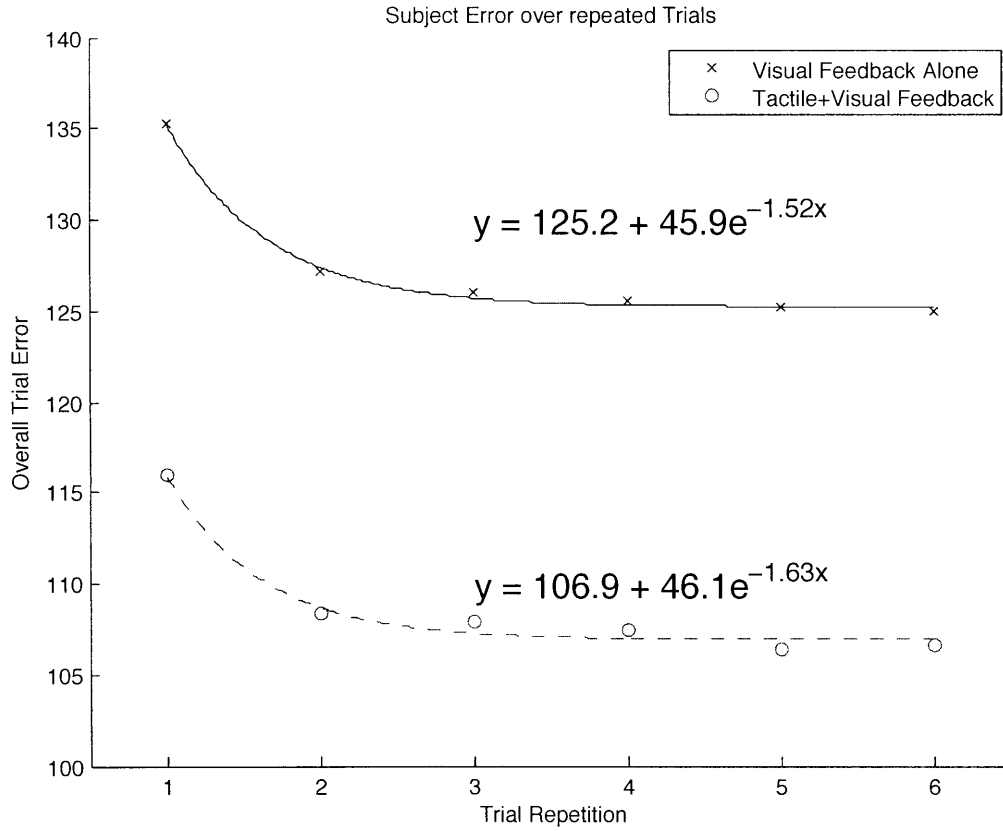
Figure 4-14: A chart of subject improvement over time, through six viewings of the same video clip. This is averaged over all 28 movies and over all 20 subjects of each type. The data points are then fit to a falling offset exponential function, indicating lower initial error, lower steady-state error, and faster settling times for the subjects using tactile feedback.

and 0.005.

Table 4.2 summarizes all of the quantitative results from this experiment.

Table 4.2: A summary of all quantitative results of the tactile feedback experiment. In dynamic cases, noticeable gains in performance were extremely statistically significant.

Test	Performance Improvement	Statistical Significance (p value)
Static positioning	5%	$p = 0.08$
Average Performance Error, all trials	21%	$p = 0.015$
Improvements over time:		
Trial 1	14.1%	$p = 0.007$
Trial 2	15.0%	$p = 0.007$
Trial 3	14.3%	$p = 0.009$
Trial 4	15.1%	$p = 0.007$
Trial 5	15.2%	$p = 0.004$
Trial 6	14.8%	$p = 0.005$
Immediate Error	11%	$p = 0.007$
Learning Time Constant	7%	$p = 0.007$
Steady State Error	15%	$p = 0.007$

Chapter 5

Conclusions and Future Work

As seen in the previous chapter, there was a noticeable change in performance with the addition of tactile feedback while learning simple motor skills. This leads us to several ideas for future improvements, applications, and possible problems with this system. We begin with those problems, leading us to possible improvements.

5.1 Problems and Improvements for Tactile Feedback Systems

The main problems with the system as it stands, are that of universal accessibility, and tuning a similar system to insure maximum performance benefits.

5.1.1 Accessibility

There are two major problems with this system, in order to make it a feasible system for the average person to employ during sports training. The first and major barrier to entry is the expense of the optical Vicon tracking system. At many thousands of dollars, it makes the wirespread use of such a system impossible. However, research in decentralized tracking systems [happening in several places concurrently] will eventually allow a system to exist completely on the user's body, with no external tracking necessary. One possible avenue of research in this field is the use of independent RF-chirping modules, that calculate relative distances in real-time to determine their layout among the body.

Secondly, the use of somewhat large vibrotactile devices will be limiting in a full-body situation, where more than 60 degrees of freedom are being monitored. This will necessitate on the order of 100 vibrotactile output devices. However, embedded clothing research is already well on the

way to developing technology to inject tactile feedback into normal clothing, either through micro-vibrotactile devices, or through small and better-tested electrotactile stimulation methods.

5.1.2 Performance

There are also several problems that need to be researched, to make this system as effective as possible.

Inaccuracies

The Vicon system, after every few minutes of testing, would either fade from completely true values over time, or would, after a temporary loss of sighting all markers, would suddenly discontinuously jump the user into a new and inaccurate tracked state. This would create sudden incorrect feedback signals for users, usually at maximum value. These signals greatly disrupt the potential of a system that is designed to subtly remove itself from the feedback loop. Any inaccuracy in tracking must be eliminated if this system is to be useful, especially for more subtle motions found in the arts such as dance, piano or other instrumental performance, or sports such as golf, where a 1 degree misalignment can make or break one's swing.

Bulk

The goal of this system is to extend to a full-body suit, regulating over 60 degrees of freedom of the human body, with over 100 actuators. At that level of parallelism, any weight of individual actuators may be distracting to certain achieved motions. This current system is on the borderline - without the control electronics and power supply, the system is extremely light. But to carry onboard power to trigger 100 of these actuators would cost roughly 50W of energy, peaking higher. This would need a reasonably massive on-board power supply such as a large battery, to sustain that for any length of time. That added weight may cause distraction for users.

That being said, the system as it stands, with power supply and control electronics off-user, is extremely portable, fluid, and lightweight, showing promise to be extended to full body at least for a proof-of-concept model.

Actuator placement and feedback design

It must be stressed that *all of the tests performed* were done without any extensive design research; that is, no extra care was taken to design an 'ideal' feedback system, or place the actuators in the ideal locations or mount them in the most comfortable and transparent manner. The feedback system was designed to be the simplest one that would accomplish the goals set out for feedback.

This was done for a specific purpose: to prove, in the worst case scenario, that this type of feedback would cause a shift in users' behavior that is marked and statistically significant. That

has been shown. Now that it has, much research needs to be done to find the proper feedback mechanisms that users will respond to most quickly and most deeply.

It is quite possible that exact actuator placement greatly influences users' ability to respond to feedback. In simple tests, movement of actuators around the upper arm greatly varied in their noticeability when signals were given. Also, proper placement can indicate what joints are to be altered, automatically, instead of requiring conscious effort on the users' part.

Also, there are many completely different styles of feedback that may be given to users when they attempt motion tasks. I applied a linear-proportional error feedback system, varying only the duty cycle of a constant 250Hz square wave. Variability in signal envelopes may present more information than just constant feedback. Variations in motor frequency and saltation frequency should show changes in performance.

Scaling to higher DOF systems

One major element that needs to be further explored is the major parallelism that will result from scaling this feedback system to a full-body model, which would comprise over 50 degrees of freedom (DOFs). In these tests we explored only up to 5 synchronous parallel signals, so true knowledge of the body's ability to interpret this massively parallel data is still not fully known. However, users rarely had issues adjusting to these levels of parallel data. No evidence so far, from this experiment or others mentioned earlier, has shown an upper bound to the bandwidth of our somatosensory system.

From the evidence gained in this experiment, it seems that there are two likely reactions to a scaled up system. First, users may completely accustom themselves to the parallel degrees of freedom. They may fully adjust over a period of hours of system use, so that they internalize the signal data and they can easily parse it into separate motor commands, in parallel. If this is not possible, a change in feedback control design could be used to always highlight and feedback only the most salient errors in motion. At this point, no long term use study has been performed; it is still my hope, however, that a user would unconsciously accustom themselves to any parallel presentation of tactile information, given adequate time to become familiar with the system. If this is not possible, switching to a feedback mechanism that chooses salient DOFs at any time to feedback is an easy modification and would still allow very similar gains in performance.

Training Regimens

It is very important, in order to get maximum benefit from this system, to properly introduce the system to new users through a series of training exercises. The genesis of this has been done in this experiment, triggering single DOF changes in static images, the simplest first step toward exploring the complex interactions with the system.

65

More techniques need to be explored to properly introduce more parallel interactions, as well as interactions with higher dynamic content. Also, not all interactions with the system need to based on real-time imitation of a teacher; it would be possible to record a motion sequence, and base the playback of the 'teaching sequence' not on real-time, but instead on the current user's performance. For instance, in a golf swing, instead of requiring the subject to perform the golf swing in real-time, instead the system may only watch the current position of the golf swing, guess which part of the real swing it most resembles, and feedback based on the current position. This would allow people to perform their motions at whatever speed they required to properly learn the motion, and they could ramp up speed as they felt confident. Similarly, learning motions at slow but fixed speeds, which ramp up based on performance, is a possibility that may allow users to learn skills very quickly, much as a musical student uses a metronome and slowly increases the tempo as they gain confidence in their motions.

The important aspect of these training regimens is that they may all be presented as options, to be chosen by the user based on their preferences.

5.2 Future Experiments and Applications

The most exciting part of this work is that is was only a proof of concept. It has shown improvements in behavior, and merits further research to see how much those benefits can be stretched. There are many possible future tests and applications worth mentioning, for future research:

5.2.1 Experiments

It would be interesting to test how quickly and accurately people could learn movements with no visual feedback whatsoever, since this is typically the primary method by which we learn motion. The test setup already allows this, by removal of the video feedback. This would test on a fundamental level the ability of tactile feedback to instruct in motions, but also would provide a test for fundamental blind training, described more below in Section 5.2.4.

Varieties of feedback methods must be tested. A simple comparison of all-joint feedback versus most-salient-joint-error feedback [where only the largest joint errors are fed back to the user] would show user's capabilities in that arena, but would be most clear when performed on a full body suit instead of an arm subset, since the smaller DOF number wouldn't allow as much difference in those feedback methods.

5.2.2 Sports and Dance

The most obvious and similar to studies already completed, is the extension of this device to the full body, for use in learning sports motions and artistic motions such as dance. The system needn't be

changed in any way, just extended, to test the usefulness of such a system.

5.2.3 Rehabilitation

One of the original goals of this system was to provide a method of easy rehabilitation for those suffering from neurological trauma such as stroke. The tests performed show promise but no direct confirmation that such a system would benefit those recovering from such a trauma. Therefore, much needs to be done to confirm that this type of system in fact does aid those in that situation.

5.2.4 Motor Learning for the Blind

A still untested but likely hypothesis is that a tactile feedback system would greatly improve the blind community's ability to learn motor tasks such as dance, which is currently a slow laborious process. No official research has been performed as of yet, so this is an untapped element of research.

5.2.5 Static Motor Learning, Posture

There is an immediate, simple, and portable application for this type of feedback - it can be used to correct posture. A much simpler tracking device could measure inclinations of a users' joints, and judge if they were in proper posture. This could all be done with very small electronics such as accelerometers. After calibration to a user, this device could provide *immediate* feedback to a user about their posture, so that they may correct it before assuming the bad posture for many minutes. It is those extra minutes that mis-train the body into holding those bad postures as habits. It is likely (but still unproven) that an immediate wearable feedback mechanism about posture, would have permanent positive effects on users' postural behavior. Furthermore this device could be made small, very inexpensive, and would likely take effect only after several hours of use, for permanent benefit.

5.3 A Final Note

This research was not intended to solve a specific problem. It was to introduce a completely new method to improve people's lives, to help them gain new abilities more effectively. The research shows great promise in this avenue. It is up to future work to utilize this new method, for everyone's benefit.

Bibliography

[1] http://en.wikipedia.org/wiki/electromyography.

[2] http://en.wikipedia.org/wiki/hm.

[3] http://vicon.com.

[4] http://www.myomo.com/technology.htm.

[5] http://www.simlog.com/index.html.

[6] R.B. Ammons. Effects of knowledge of performance: a survey and tentative theoretical formulation. *Journal of General Psychology*, pages 279–99, 1956.

[7] E. Bilodeau and I Bilodeau. Some effects of introducing and withdrawing knowledge of results early and late in practice. *Journal of Experimental Psychology*, 58:142–44, 1959.

[8] E. Bilodeau and I Bilodeau. Motor-skills learning. *Annual Review of Psychology*, 12:243–280, 1961.

[9] E.A. Bilodeau. Acquisition of two lever-positioning responses practiced over several periods of alternation. *Journal of Experimental Psychology*, pages 43–49, 1953.

[10] E Bizzi, F.A. Mussa-Ivaldi, and R. Shadmehr. System for human trajectory learning in virtual environments, 1996.

[11] B Brooks, J.E. McNeil, and F.D. Rose et al. Route learning in a case of amnesia: a preliminary investigation into the efficacy of training in a virtual environment. *Neuropsychological Rehabilitation*, 9:63–76, 1999.

[12] J.S. Brown. A proposed program of research on psychological feedback (knowledge of results) in the performance of psychomotor tasks. In *Research Planning Conference on Perceptual and Motor Skills. AF HRRC Conf. Report 49-2*, pages 81–87, 1949.

[13] U.S. Census Bureau. Disability status: 2000. 2000.

[14] R.W. Cholewiak and A.A. Collins. The generation of vibrotactile patterns on a linear array: Influences of body site, time and presentation mode. *Psychophysics*, pages 1220–1235, 2000.

[15] J.E. Conklin. Linearity of the tracking performance function. *Perceptual Motor Skills*, 9:387–91, 1959.

[16] S. Corkin. Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, 6:255–265, 1968.

[17] J.C. Craig and C.E. Sherrick. *Tactual perception: A source book*, pages 209–233. Cambridge University Press, New York, 1982.

[18] Ian R. Summers Ed. *Tactile Aids for the Hearing Impaired*. Whurr Publishers, London, 1992.

[19] F.A. Geldard and C.E. Sherrick. The Cutaneous "Rabbit": A Perceptual Illusion. *Science*, 178:178–179, October 1972.

[20] M. Goldstein and C.H. Rittenhouse. Knowledge of results in the acquisition and transfer of a gunnery skill. *Journal of Experimental Psychology*, pages 187–96, 1954.

[21] Goodwin, McCloskey, and Matthews. The contribution of muscle afferents to kinaesthesia shown by vibration induces illusions of movement and by the effects of paralysing joint afferents. *Brain*, 95:705–748, 1972.

[22] E. Gunther, G. Davenport, and S. O'Modhrain. Cutaneous grooves: Composing for the sense of touch. In *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02), Dublin, Ireland*, pages 1–7, 2002.

[23] M Hancher. A motor control framework for many-axis interactive robots. Master's thesis, 2003.

[24] M Holden. Virtual environments for motor rehabilitation: Review. *CyberPsychology and Behavior*, 8(3), 2005.

[25] C. Wall III and M.S. Weinberg. Balance prostheses for postural control. *IEEE Engineering in Medicine and Biology Magazine*, pages 84–90, March/April 2003.

[26] L. Jones. Development of a tactile vest. In *Proceedings of the 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 1–2, 2004.

[27] K.A. Kaczmarek, J.G. Webster. P. Bach y Rita, and W.J. Tompkins. Electrotactile and vibrotactile displays for sensory substitution systems. *IEEE Transactions on Biomedical Engineering*, 38(1):1–16, 1991.

[28] H. Kajimoto, N. Kawakami, S. Tachi. and M. Inami. Smarttouch [haptic interface]. *Computer Graphics and Applications IEEE*. 24(1), 2004.

[29] J. Lieberman. Teaching a robot manipulation skills by demonstration. Master's thesis, Massachusetts Institute of Technology, 2002.

[30] FD Rose, E.A. Attree, and B.M. Brooks. Training in virtual environments: transfer to real world tasks and equivalence to real task training. *Ergonomics*, 43:494–511.

[31] M. Rothenberg, R.T. Verrillo, S.A. Zahorian, M.L. Brachman, and S.F. Bolanowski Jr. Vibrotactile frequency for encoding a speech parameter. *Journal of the Acoustical Society of America*, (4):1003–1012, 1977.

[32] A.H. Rupert. An instrumentation solution for reducing spatial disorientation mishaps: a more 'natural' approach to maintaining spatial orientation. *IEEE Engineering in Medicine and Biology*, pages 71–80, 2000.

[33] C. Sherrick. A scale for rate of tactual vibration. *JASA*, 1(78), 1985.

[34] A.F. Smode. Learning and performance in a tracking task under two levels of achievement information feedback. *Journal of Experimental Psychology*, pages 297–304, 1958.

[35] F.V. Taylor. Simplifying the controller's task through display quickening. *Occupational Psychology*, pages 120–25, 1957.

[36] E Todorov, R Shadmer, and E Bizzi. Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task. *Journal of Motor Behavior*, 29:147–158, 1997.

[37] J.M. Weisenberger. Sensitivity to amplitude modulated vibrotactile signals. *Journal of the Acoustical Society of America*, 80:1707–1715, 1986.

[38] J. Zwislocki, R.P. Hellmann, and R.T. Verrillo. Threshold of audibility for short pulses. *Journal of the Acoustical Society of America*, 34:1648–1652, 1962.