REAL-TIME SOMATOSENSORY FEEDBACK FOR NEURAL PROSTHESIS CONTROL: SYSTEM DEVELOPMENT AND EXPERIMENTAL VALIDATION

by

Daniel H. Bacher

Bachelor of Science in Biomedical Engineering, Syracuse University, 2007

Submitted to the Graduate Faculty of

Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Master of Science in Bioengineering

University of Pittsburgh

2009

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This thesis was presented

by

Daniel H. Bacher

It was defended on

July 16, 2009

and approved by

Roberta Klatzky, Professor, Department of Psychology, Carnegie Mellon University

George Stetten, Professor, Department of Bioengineering

Douglas Weber, Assistant Professor, Department of Bioengineering

Thesis Advisor: Aaron Batista, Assistant Professor, Department of Bioengineering

REAL-TIME SOMATOSENSORY FEEDBACK FOR NEURAL PROSTHESIS CONTROL: SYSTEM DEVELOPMENT AND EXPERIMENTAL VALIDATION

Daniel H. Bacher, M.S.

University of Pittsburgh, 2009

Recent advances in neural prosthetics have provided patients with the ability to use signals derived from motor areas of the cerebral cortex to directly control an external device under visually guided closed-loop control. To attain a more natural form of prosthesis control, it is desirable to develop systems capable of providing real-time somatosensory feedback as well as visual feedback, akin to how we naturally process sensory information to control our limbs. To this end, a sophisticated data acquisition, control and feedback system was developed for neural prosthetics and psychophysics research. The system deterministically collects and processes high volume neural ensemble activity, limb kinematics, and eye movements while generating visual stimuli in an immersive three-dimensional virtual reality (VR) environment. A vibrotactile feedback device was also developed and incorporated into the system. It delivers real-time limb kinematics feedback in the form of continuous, graded vibratory stimulation. A flexible and intuitive user interface allows the researcher to design experimental paradigms and adjust parameters on the fly during experiments.

A psychophysical study was conducted using this system to evaluate the potential use of vibrotactile feedback as a sensory substitution method to provide somatosensory feedback for neural prosthesis control. The study also aimed to provide insight into the mechanisms of multimodal sensory processing and sensory-motor control. Able-bodied human subjects performed a trajectory-following reach task in the VR environment under different visual and

vibrotactile feedback conditions. The study showed that vibrotactile feedback is capable of enhancing motor performance, implying that subjects were able to integrate and effectively use this new 'proprioceptive-like' sensory modality. Subjects were also able to partially maintain task performance using vibrotactile feedback in the absence of visual feedback. Improved motor learning and motor skill consolidation were also observed after training in the VR environment with vibrotactile feedback. These results suggest that vibrotactile feedback may be a viable method for delivering somatosensory feedback for applications such as neural prosthesis control, motor rehabilitation, and enhanced human-computer interaction.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0   INTRODUCTION

Despite recent advances in modern medicine, there exists an ever-increasing need for viable clinical treatments for people suffering from spinal cord injury, limb loss and other neurological disorders that leave patients "locked-in", unable to communicate or interact with their environment. In the United States alone there are 1.2 to 1.5 million people living without a limb, with 185,000 amputations performed each year (ACA 2005). There are approximately 250,000 people with spinal cord injuries in the United States (Center 2005). Moreover, an estimated 30,000 Americans are suffering from ALS, a neurodegenerative disease that leads to paralysis (ALSA 2007).

Research in the field of Brain-Machine Interfaces (BMIs), or more specifically neural prosthetics, aims to create a direct link between the nervous system and machine with hopes of one day developing advanced medical devices capable of restoring motor function in those suffering from neurological disorders and nervous system injury. Building on years of scientific studies investigating the role of sensory and motor cortical areas of the brain in the control of movement, numerous laboratories are working toward developing closed-loop BMIs, that is to create an interface allowing a subject or patient to control a computer cursor, robotic arm or some other external device solely using their movement intention or 'thoughts', and receive feedback from the device in real-time.

In parallel with progress made in neural prosthetics there have been significant advances in the field of Human-Computer Interaction (HCI). Recently developed technologies are revolutionizing how we interact with our computers and electronic devices, including three-dimensional (3D) virtual reality displays and haptic feedback devices capable of allowing users to physically interact with objects in virtual environments. Many of these devices use the technique of sensory substitution- providing feedback to a user through a sensory channel not normally used to encode that type of stimulus. One particular example of such a sensory substitution system in the HCI field is the use of haptic feedback to encode visual stimuli in assistive devices for the blind. Such haptic feedback devices, or devices that apply force and/or vibratory stimuli to the skin, have recently been considered for various other applications, including motor rehabilitation, and most recently, BMI applications.

This thesis describes work towards developing a sophisticated data acquisition, control and feedback system to be used in neural prosthetics research and sensory-motor psychophysics. This system is comprised of data acquisition hardware and software, custom experimental control software, an immersive 3D virtual reality display, and a vibrotactile feedback device. This work also describes a human psychophysical study using this system, which investigates the potential use of vibrotactile feedback for a variety of applications, including neural prosthesis control, motor rehabilitation, and enhanced human-computer interaction.

2.0   BACKGROUND

2.1   BRAIN-MACHINE INTERFACES

"Brain-Machine Interface" (BMI) is a term commonly used to describe any device that interfaces the central nervous system to an electronic device. The term Brain-Computer Interface (BCI) is also commonly used, as most often the nervous system is linked to a computer or collection of computers. There are various methods (non-invasive or invasive) to record signals from the cenral nervous system. Non-invasive techniques such as electroencephalography (EEG) and magnetoencephalograhy (MEG) indirectly measure the electrical and magnetic fields respectively generated by the brain. Both of these methods have various other clinical uses, but recently have been used for BMI applications. Electrocorticography (ECoG) is an invasive technique where a grid of electrodes is placed on the surface of the cerebral cortex to record neural activity. Implanting these ECoG grids requires the removal of the portion of the skull, but is commonly used in detecting seizure loci in patients with epilepsy. An even more invasive technique, but one that has been extensively used in BMI research, is implanting either single or multiple electrodes directly into the cerebral cortex. This method requires penetrating electrodes into the cortical tissue in addition to making a hole in the skull to implant the electrodes, but unlike ECoG, single electrode recordings capture the activity of single neurons in the brain.

Schwartz et al. (2006) provides a detailed overview of current neurophysiological recording technology.

2.1.1   Neuromotor Prostheses

A neuromotor prosthesis, or simply neural prosthesis, is a device that directly interfaces with the nervous system to restore motor function in patients suffering from spinal cord injury, limb loss and other neurological disorders resulting in loss of voluntary movement control. Recent work in the field of neural prosthetics has demonstrated the ability of non-human primates (Carmena et al. 2003, Santhanam et al. 2006, Taylor et al. 2002) and humans (Hochberg et al. 2006) to control a neural prosthesis under visually guided closed-loop control. In these studies, microelectrode arrays (a chip containing up to 100 electrodes) implanted in motor areas of the cerebral cortex simultaneously record the activity of dozens to hundreds of neurons. Computer algorithms then interpret the pattern of neural firing rates to generate output control signals to an external device solely based on the user's movement intent, allowing the user to control the device just by 'thinking' about moving it.

The most common configuration of a neural prosthesis in non-human primate research involves tracking arm and hand position (and sometimes also eye position) in two or three dimensions while recording neural ensemble activity as the subject performs some form of reach task (Schwartz et al. 2006). A statistical model, or decoder, can then be trained to predict movement kinematics based on neural ensemble activity. These decoded signals are then used to control a cursor on a computer screen, or more recently a robotic arm (Velliste et al. 2008).

2.1.2 The Components of a Neural Prosthesis System

A typical neural prosthesis can be decomposed into three functional categories: data acquisition sub-systems, a central control unit, and output devices. The data acquisition sub-systems are used to record and process various measurements from a subject, including neural ensemble activity, arm/hand kinematics and eye movements. Multiple companies (Tucker Davis Technologies Inc., Alachua, FL; Cyberkinetics Inc., Foxboro, MA; Plexon Inc., Dallas, TX; Ripple LLC, Salt Lake City, UT) have developed hardware and software that is capable of acquiring and processing high volume neural data, typically from microelectrode arrays (Blackrock Microsystems Inc., Salt Lake City, UT). Two or three dimensional arm kinematics data are obtained either through the use of optical motion capture systems (e.g. Northern Digital Inc., Waterloo, Ontario, Canada; Vicon Inc., Los Angeles, CA; Phasespace Inc., San Leandro, CA) or robotic manipulandums (e.g. PCB Piezoelectronics Inc., Depew, NY). Eye position is measured using infrared optical systems (e.g. SR Research Ltd., Ottawa, Ontario, Canada) or magnetic eye coil systems (Robinson 1963).

All of these data (neural, arm movement, eye movement) are piped into the central control unit. This central control unit dictates experimental task flow, aligns and processes the streaming input data, and generates output commands based on the input data and the particular behavioral control paradigm. Numerous third party software packages have been developed to function as a central control unit (e.g. Bryant & Gandhi 2005; Ojakangas et al. 2006; Tempo by Reflective Computing Inc., WA). Many laboratories also develop software in-house for custom applications.

Lastly, output devices provide prosthesis state information (e.g. position, velocity) to the user. Typically this is done by rendering a cursor on a computer screen representing hand

position either on a standard 2D screen (Hochberg et al. 2006) or in a 3D virtual reality environment (Taylor et al. 2002). Another type of controllable output device is an anthropomorphic robotic arm, which was recently used by non-human primate subjects to feed themselves under neural control (Velliste et al. 2008). In addition, other output devices can be used to present non-visual forms of feedback, such as tactile and auditory displays. These non-visual displays have been developed for a wide array of applications and will be discussed in subsequent sections. One promising method for providing non-visual feedback for neural prosthetics applications is through direct cortical microstimulation (London et al. 2008). Cortical microstimulation has been extensively researched, but has yet to be effectively implemented as a means of providing dynamic feedback for neural prosthesis control. Each of these system output methods attempts to 'close the loop' by conveying informative sensory feedback to allow for natural control of the prosthesis.

2.2   SOMATOSENSORY FEEDBACK

2.2.1   Sensory Substitution

The term sensory substitution can be defined as "…the use of one human sense to receive information normally received by another sense" (Kaczmarek et al. 1991). Sensory substitution has been applied in the clinical setting to supplement or replace a deficient or lost sensory modality. One classic example of this approach was the development of a tactile-vision sensory

substitution system (Bach-y-Rita et al. 1969), which was comprised of a vibrotactile display that encoded a visual scene captured by a video camera. Another more recently developed system is a balance prosthesis (Vuillerme et al. 2007) that uses electrotactile stimulation of the roof of the mouth to encode a pressure map representing the force distribution on the user's feet to correct posture in those suffering from balance disorders.

Tactile stimulation of the skin is a common choice for sensory substitution systems due to the ease of creating devices capable of eliciting tactile perception. Tactile perception also provides an available sensory channel during most activities. Although we are constantly using our sense of touch to interact with our environment, tactile stimulation is capable of providing rich sensory signals while avoiding interference with other sensory channels such as visual or auditory processing.

2.2.2   Tactile Perception

The skin contains many different types of tactile receptors, whose properties have been well characterized and described in the literature (e.g. Kaczmerak et al. 1991). Vibrotactile stimulation activates a particular subset of these receptors, namely pacinian corpuscles and meissner's corpuscles. Pacinian corpuscles are quickly adapting mechanoreceptors which respond to vibratory stimuli with a frequency range of approximately 40 to 800 Hz with a peak sensitive of around 200-300 Hz (Kaczmerak et al. 1991). Meissner's corpuscles are also quickly adapting receptors that respond to vibration and flutter (stimulation below ~40 Hz). The responsive frequency range for Meissner's corpuscles is approximately 10 to 200 Hz with peak sensitivity around 20 to 40 Hz (Kaczmerak et al. 1991). Other mechanoreceptors in the skin

respond to different forms of tactile stimuli, including ruffini endings which are sensitive to stretch, shear, and tension, and merkel disks, which respond to edge pressure (Kaczmerak et al. 1991).

The density of these mechanoreceptors varies for different areas of the skin on the body. This factor plays a critical role in the development of tactile displays. For example, fingertips contain a very high density of receptors, allowing for very fine discrimination of multiple tactile stimuli. The back, however, has a more sparse concentration of receptors, which means that if a dense array of tactile actuators were applied to the back, a person would not be able to discriminate between the actuators that are very close to each other. The high tactile acuity of the fingertips also makes them an ideal location to deliver highly precise sensory information, but depending on the application covering the fingertips could impede the natural use of the hand.

2.2.3 Tactile Sensory Substitution Systems

Recent advances in commercial hardware have made it possible to build low-cost tactile displays, making tactile sensory substitution devices appealing for a variety of commercial and medical applications. Numerous studies have been conducted investigating the feasibility of using tactile displays as an effective sensory substitution method. One recent study (Marston et al. 2007) evaluated the use of simple tactile and auditory displays to guide blind subjects along an instructed path. Results of the study concluded that even simple binary tactile cues were capable of effectively directing the subjects along the correct path.

Other systems have been developed for more broad applications including motor rehabilitation, sports training and enhanced human-computer interaction. One such system, coined TIKL: Tactile Interaction for Kinesthetic Learning (Lieberman & Breazeal 2007), provides vibrotactile feedback of joint angles integrated into a wearable robotic device. This system aims to enhance the awareness of the limb in space to help train new motor skills or even recover motor skills lost due to neurological disease or injury.

Many experimental systems have been created that all rely on using tactile feedback to encode some other form of information to the user of the device. Other examples include an auditory-tactile system for the hearing impaired (Reed and Delhorne 1995), a magnetic compass-tactile system (Nagel et al. 2005) that provides greater directional awareness, and a body orientation-tactile system (van Erp and van Veen 2003) for astronauts to be used in zero gravity environments.

2.3   VIBROTACTILE FEEDBACK FOR BRAIN-MACHINE INTERFACES

Natural control of our limbs is coordinated not only by visual feedback but also somatosensory feedback. The awareness of our limbs in space, referred to as proprioception, helps coordinate limb movements in addition to allowing us to make movements without having to fixate on the limb. As described earlier, BMIs have been used under visually guided closed-loop control, but there remains a need for a system to provide somatosensory feedback to provide more natural BMI control.

One promising method for providing this non-visual feedback is through vibrotactile feedback. Using sensory substitution, vibrotactile feedback could encode the missing proprioceptive feedback to better control the neural prosthesis. Vibrotactile feedback could also be used to encode the tactile sensation of the prosthesis. For example, if a BMI user were controlling a robotic arm, vibrotactile feedback applied somewhere perceptible to the subject could relay a signal when the robot bumps into or slides past an object. One other advantage of vibrotactile feedback is that it provides a much faster form of feedback than visual feedback. Visual perception is a relatively slow process compared to tactile perception, so by reintroducing a form of feedback that has comparable temporal characteristics, a more sophisticated form of control might be attained.

Two recent studies investigated the use of vibrotactile feedback for Brain-Computer Interface (BCI) control. In contrast to the neuromotor prosthesis described earlier, these studies used electroencephalography (EEG) to measure neural activity of human subjects. The modulation of the EEG signal was then decoded and used to control the prosthesis. In both studies (Chatterjee et al. 2007 and Cincotti et al. 2007), subjects demonstrated one-dimensional control of a computer cursor with either visual or vibrotactile feedback. Both studies show that vibrotactile feedback is sufficient to effectively control the prosthesis in these experimental tasks. Further studies are needed to determine how effective vibrotactile feedback could be for more sophisticated BMI control.

3.0   SYSTEM DEVELOPMENT

Based on the general architecture described in Section 2.1.2, a complete data acquisition, experimental control, and feedback system was developed for non-human primate neural prosthetics and human psychophysics research. The system is comprised of a combination of third party hardware and software for data acquisition, a custom LabVIEW-based (National Instruments Inc., Austin, TX) software package for system control, and a 3D virtual reality (VR) display and vibrotactile feedback device. Data acquisition sub-systems record and transmit neural ensemble data, limb position, eye position, and other general-purpose data to a central control unit consisting of real-time (RT) controller and Host PCs. The central control unit synchronizes and processes incoming data and generates the appropriate output commands to the VR and vibrotactile displays as dictated by the experimental paradigm design.

3.1   SYSTEM DESIGN

A detailed system architecture diagram is shown in Figure 1 below.
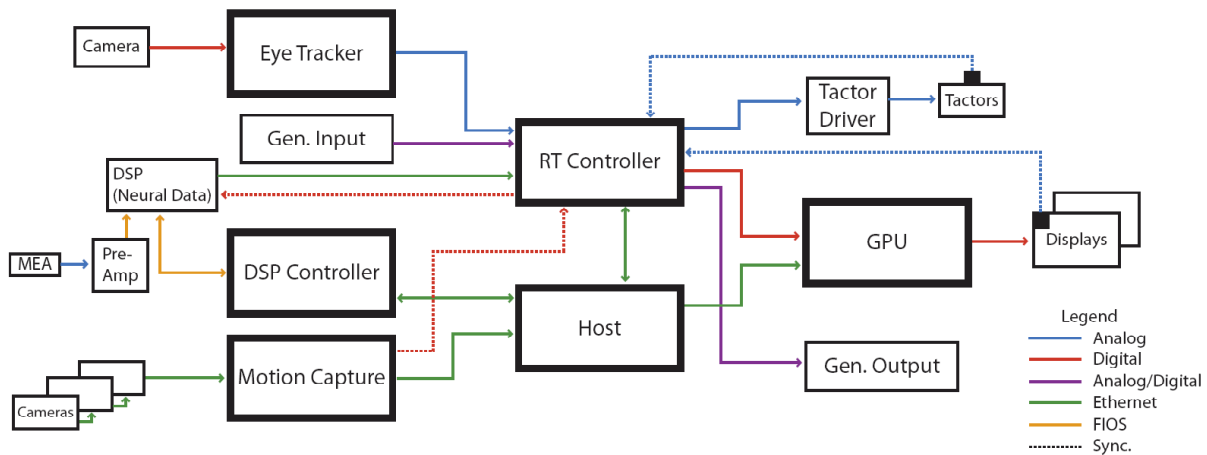
Figure 3.1: System Architecture. Data flows from left to right, with data acquisitions sub-systems on the left (eye position, neural data, arm position), the central control unit in the middle (RT controller and Host) and output devices on the right (graphics display, tactile feedback system). Abbreviations: MEA = microelectrode array (to directly record single unit neural activity), DSP = digital signal processor (to acquire and process neural data), Gen. Input / Gen Output = general-purpose input/output, GPU = graphical processing unit. Data communication: each communication line is color coded- blue = analog, red = digital, purple = both analog and digital, green = ethernet (TCP/IP or UDP), and orange = fiber optics (FIOS). Dashed lines denote synchronization signals. Each box with a thick border represents a separate networked PC.

3.1.1   Data Acquisition Sub-systems

For neural prosthetics studies, neural activity, arm position and eye position can all be simultaneously recorded. For human psychophysics, it may only be necessary to record arm and hand movements, and possibly eye movements depending on the study. The following is a summary of the data acquisition hardware and software.

3.1.1.1 Neural Ensemble Recordings

Neural data is obtained using Tucker Davis Technologies Inc's (TDT) neurophysiology digital signal processing system. The TDT pre-amplifier interfaces directly with the head-stage of a 96-electrode array (Blackrock Microsystems Inc.) implanted in a motor area of a non-human primate's cerebral cortex. The raw spiking waveforms are amplified, digitized and transmitting via fiber optics to the RZ-2 digital signal processor (DSP), where real-time filtering and spike-sorting is performed. Utilizing a new Ethernet (UDP) interface, real-time neural spiking events are sent to the RT controller every millisecond (sampling rate = 1 kHz). A DSP controller PC is used to program the RZ-2 processor through a fiber-optic interface. The DSP controller also allows the researcher to visualize raw waveforms and individual sorted spikes, while providing the flexibility to adjust filter and sorting parameters online. Raw data and all detected action potential waveforms with time stamps are stored on the DSP controller as well for offline analysis.

3.1.1.2 Motion Capture

Movement kinematics of limb position are measured using Phasespace Inc.'s active-marker optical motion capture system. Up to dozens of markers (red/infrared LEDs) can be simultaneously tracked in three dimensions by aiming a series of networked cameras at the workspace. LEDs are individually strobed at a high rate by a wireless controller to uniquely identify each marker. After a brief calibration procedure, the system is capable of tracking the markers with millimeter accuracy at sampling rates up to 480 Hz. The position (x,y,z coordinates), frame number and identity of each tracked marker are sent to the Host from the Phasespace Base Station ('Motion Capture' PC in Figure 3.1) via an Ethernet (TCP/IP) connection.

3.1.1.3 Eye Tracking

Three-dimensional eye tracking is achieved using SR Research Ltd.'s infrared (IR) binocular eye tracking system and custom written LabVIEW software for calibration. Figure 3.2 shows the geometry of the eye tracker setup: A single camera is mounted vertically above the subject (green rectangle) and using a mirror mounted at 45 degrees relative to the floor, obtains a clear view of both of the subject's eyes through a pair of 'cold' mirrors. Because the cold mirrors allow IR energy to pass through while reflecting the visual spectrum, the IR camera can effectively track the pupils and corneal reflections of both eyes through the mirrors. The camera-mirror setup allows for a large workspace for the subject to move their arm, as opposed to mounting the camera directly behind the mirrors. An IR illuminator is also placed in parallel with the camera to provide IR contrast for the pupil and corneal signal. For binocular tracking, the system's sampling rate is 500 Hz. Raw, un-calibrated x and y positional data from each eye are sent to the RT controller as analog voltages. A custom calibration procedure is then used within the LabVIEW software on the RT controller to track eye position in 3D: The subject fixates on a series of targets presented in the 3D workspace as the software develops a generalized linear model which maps the four inputs (x and y voltages from both eyes) into 3D coordinates in the VR environment.
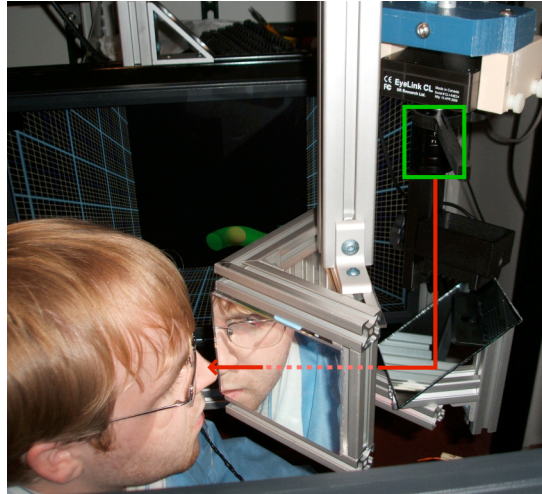
Figure 3.2: Eye tracker setup

3.1.1.4 General Purpose Input

The system is also capable of acquiring up to 24 additional general-purpose analog or digital inputs at a rate of 1 kHz. Examples include multiple EMG channel recordings and push-button inputs. A subset of these inputs channels were crucial for synchronization and timing purposes as well, which will be discussed later.

3.1.2 Central Control Unit

3.1.2.1 RT Controller

The core of the entire system is the Real-Time (RT) controller, which is an off-the-shelf PC (Intel Q6600 Quad-core CPU, 4 GB RAM) running a real-time operating system (Phar-Lap ETS, Intervalzero Inc.). The RT controller maintains a two-way communication line (Ethernet, TCP/IP) with the Host computer running LabVIEW-based software in a Windows environment. Building off of a software package created by Bryant and Gandhi (2005) used in oculomotor

research, custom LabVIEW software was developed capable of real-time, deterministic experimental control for neural prosthetics and psychophysics applications. The RT controller operates on a 1 kHz clock, meaning the system deterministically reads in, processes, and generates outputs every millisecond. Hardware from National Instruments (NI PCI-6229 Multifunction DAQ & NI PCI-6273 Analog Output Card) interfaces with the RT controller for data acquisition, inter-system communication, and generating output signals. On the RT controller, experimental task flow is structured as a finite state machine, in which each trial event is dictated by transitions from state to state as designed by the researcher. All time-critical operations are executed by the RT controller, including acquisition and alignment of incoming data, processing of data and trial timing events, maintaining system synchronization, and generating commands to output systems. Real-time neural prosthetic decoding algorithms are also executed on the RT controller.

3.1.2.2 Host CPU

The 'Host' is also a standard off-the-shelf PC (Intel Q6600 Quad-core CPU, 4 GB RAM) and it runs the LabVIEW-based graphical user interface (GUI) for the researcher in a windows environment. Through a series of simple, intuitive windows (see Figure 3.3 below), the researcher is able to easily design behavioral control paradigms by constructing and modifying state tables. Other parameters are set through a series of pop-up windows. Stimulus timing, location, and appearance can all be dynamically adjusted on the fly during experiments through the GUI. The GUI also provides real-time plots of behavioral data, including hand position, eye position, in addition to displaying state parameters, and other trial information. The Host also serves as a waypoint for the motion capture data before being sent to the RT controller due to

16

incompatibilities of the Phasespace system with the real-time operating system on the RT

controller. Software on both the RT controller and Host were compiled to executables, allowing

for stand-alone installation and execution of the software without needing to install LabVIEW.
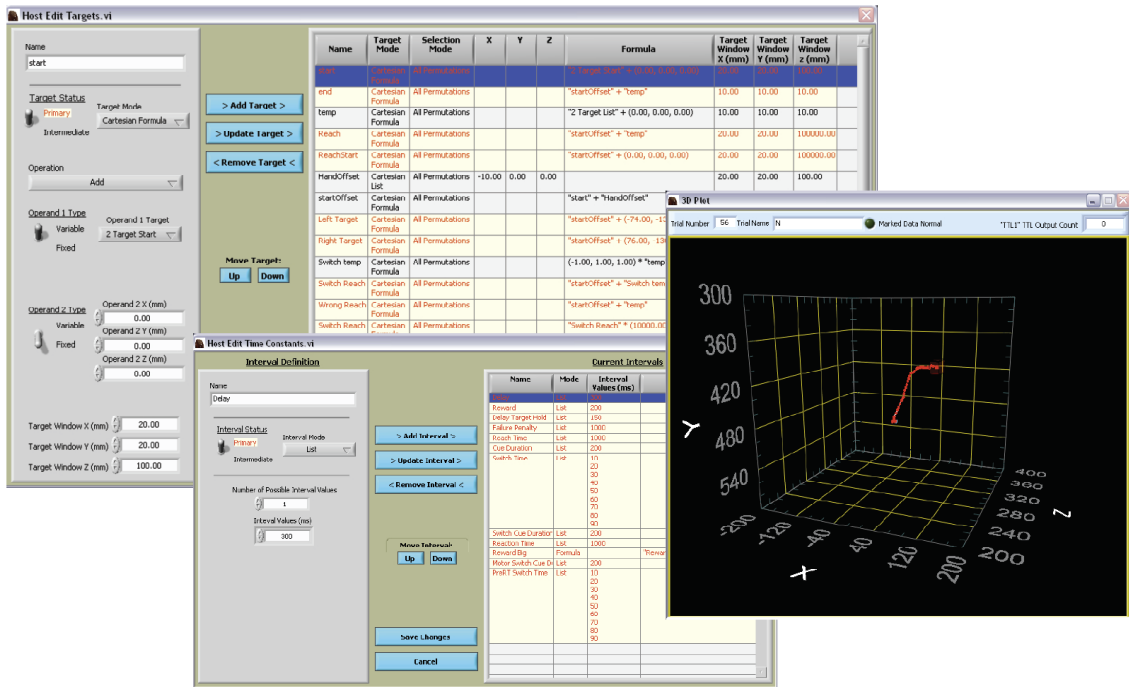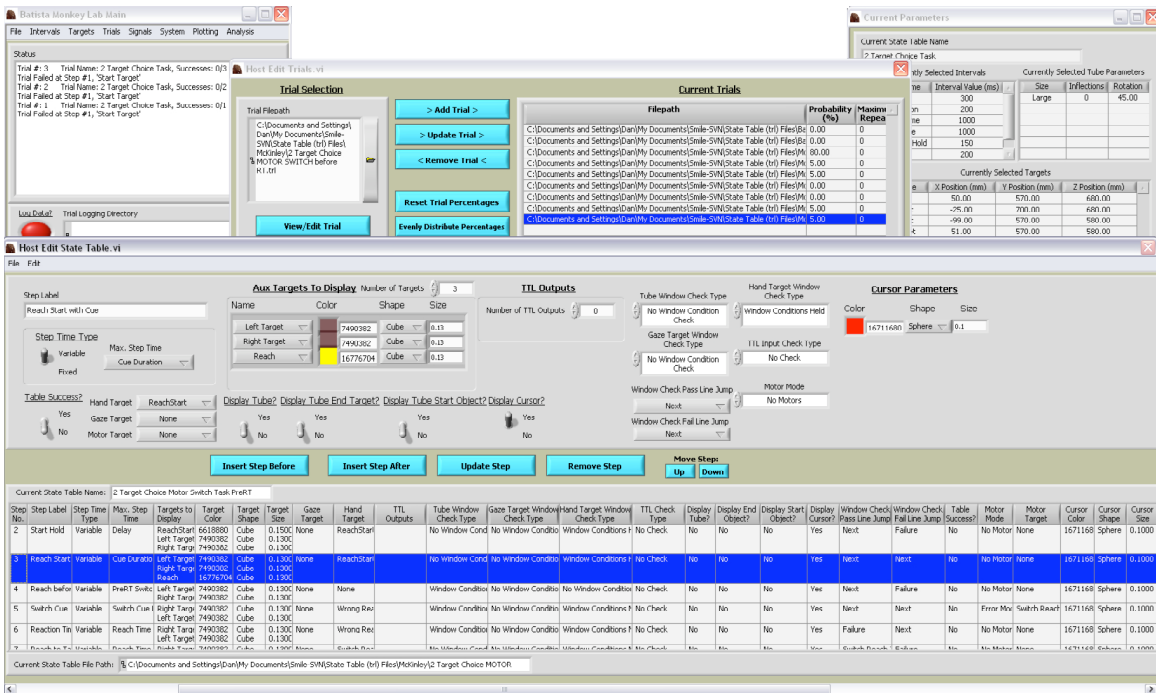
Figure 3.3: Screen captures of the researcher graphical user interface. Top panels show trial status, current parameters, trial selection and trial state table windows. The 'Host Edit State Table' window shows an example behavioral paradigm. Bottom panels show windows for specifying target location (i.e. visual or tactile stimuli) and timing intervals (dictating precise timing of trial events). The Bottom right panel shows a 3D plot of an example hand trajectory in the 3D workspace.

3.1.3 Output Devices

3.1.3.1 VR Display

Visual stimuli are presented in a 3D virtual-reality environment using a Wheatstone stereoscope configuration (Wheatstone 1838, Wheatstone 1852) as shown in Figure 3.4. A dedicated graphics computer (GPU), running in a Windows environment, receives information from the Host via Ethernet containing stimulus parameters and from the RT controller an 8-bit digital word encoding precise state transition timing information. During execution, the Host pre-loads all of the necessary information for an upcoming trial, and the RT Controller controls the precise timing of the graphical events. The GPU then renders the visual stimuli by driving two LCD monitors that are positioned on both sides of the subject's head. Two mirrors oriented at 45-degree angles relative to the LCD monitors reflect the images from the monitors to each eye independently (see Figure 3.4), allowing for true stereoscopic perception of the virtual scene. A third monitor is shown in Figure 3.4, which is used as a standard 2D display for non-human primate studies.
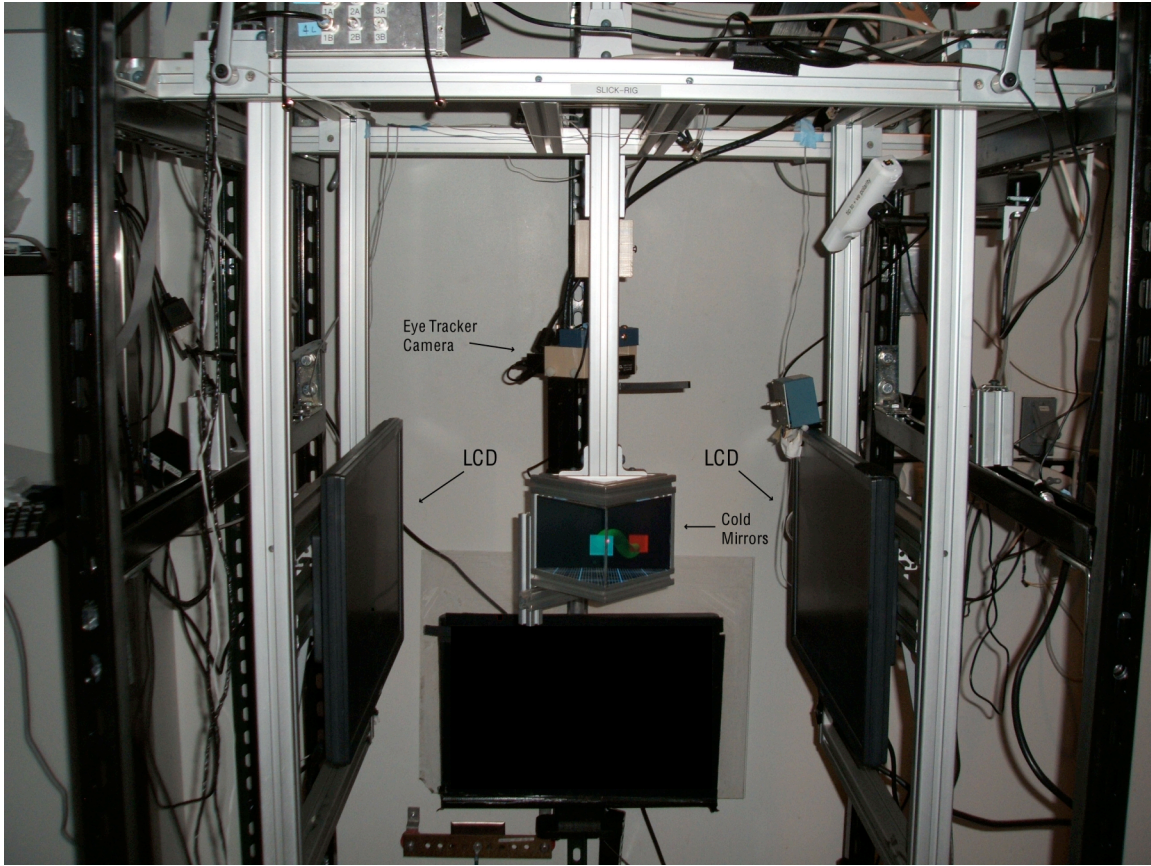
Figure 3.4: Stereoscopic VR display setup. Two LCD monitors project onto a pair of mirrors that display independent images to each eye. Subjects place their nose at the median of the mirror setup as shown in Figure 3.1. A visual disparity is added to the display to mimic natural binocular vision.

Graphics were generated with LabVIEW's OpenGL perspective drawing software package. Using the flexibility of this package, a wide variety of 3D objects can be rendered in the VR environment, complete with programmatically controlled effects and depth cues, including shadowing effects, relative size disparity, object occlusions, and variable transparencies. To emulate stereoscopic vision of the VR scene, a disparity between the two 'virtual cameras', or virtual vantage points, was introduced to mimic the natural disparity between the eyes. A careful calibration procedure was conducted which mapped the virtual coordinate system of the software into the physical coordinate axes as defined by the motion

capture coordinate system: this was accomplished by placing multiple markers on a 3D object in the workspace and applying the correct transformation to the arbitrary LabVIEW VR coordinate system to make the rendered marker positions precisely match the true motion capture marker locations. This calibration allows the researcher to define the location of visual stimuli in accurate physical coordinates in the 3D workspace. Figure 3.5 below shows an example screen capture from the VR display for a trajectory-following reach task.

Figure 3.5: Example image from one of the monitors in the VR display. In this example, a subject is performing a trajectory-following reach task where they reach from the blue cube to the red cube along the green tube. Hand position is rendered as the light red sphere in the 3D workspace. The purple 'x' is used as a fixation target to control gaze.

3.1.3.2 Vibrotactile Feedback System: VIPER

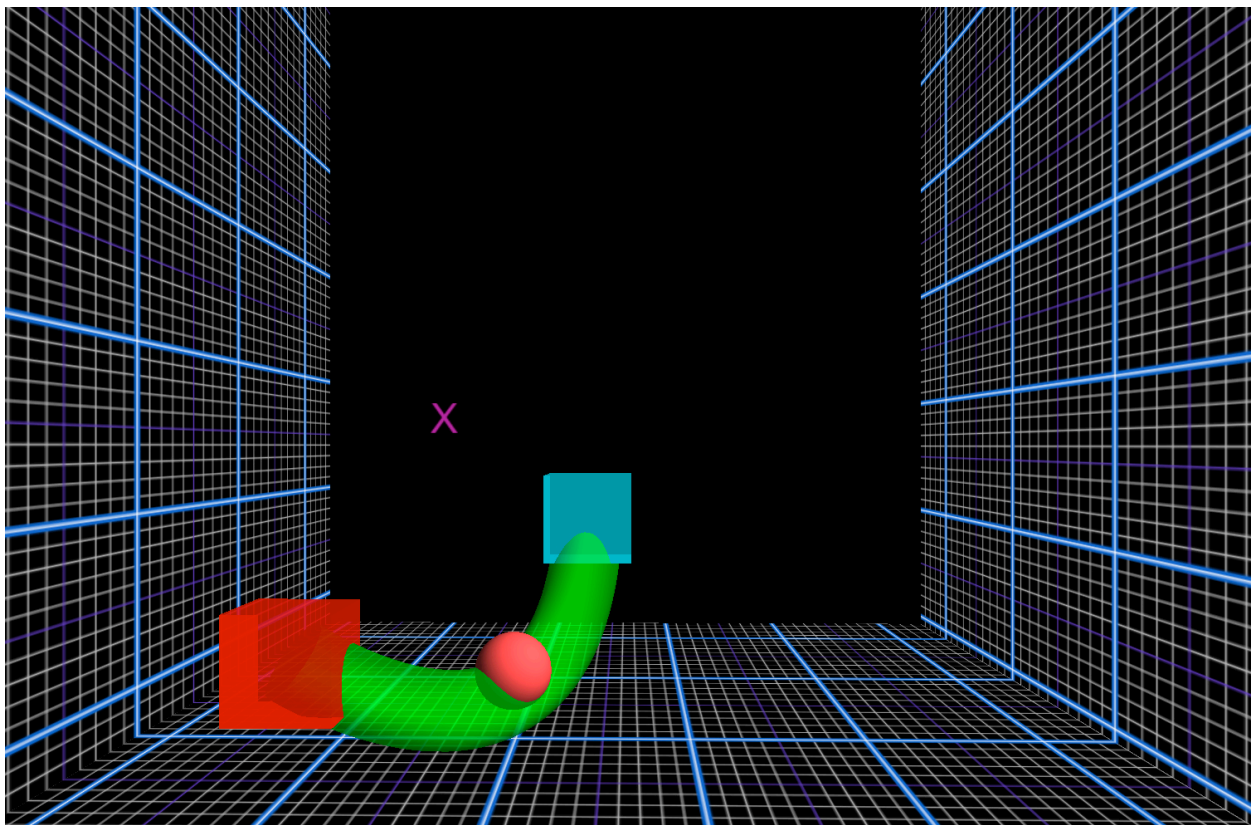A custom-built vibrotactile feedback system, coined VIPER (Vibrotactile Proprioceptive EmulatoR), provides real-time arm state information to the subject through vibratory stimulation delivered by small DC pager motors (Marlin P. Jones & Assoc. Inc., Lake Park, FL) and custom hardware. The motors are driven with analog voltages from the RT controller, which are current buffered to provide sufficient power. With an off-centered, rotating mass, the analog voltage spins the unbalanced motor, resulting in a vibration at the frequency of the rate of rotation. For a voltage range between 0.5 and 3.0 volts, the motors rotate from approximately 55 to 110 Hz. Graded vibrations within this range can be applied, with the voltage sent to the motor being updated every millisecond by the RT controller. Below 0.5 volts there is not enough power to turn the motor on and above 3 volts the frequency saturates near the motor's maximum output capacity. A detailed characterization of the same type of tactor (i.e. motor) is also described in Jones et al. (2005).

The current VIPER system allows for up to six actuators to simultaneously (or with any desired temporal pattern) deliver tactile stimuli to wherever the tactors are placed on the subject's body. The frequency of each tactor can be independently controlled to provide continuous, graded vibratory stimulation as opposed to just a binary on/off stimulus. Custom plastic casings with an outer layer of Velcro were built for the motors to make them more durable and provide an easy method of attaching them to the subject. One simple implementation of a two-tactor feedback glove is shown below in Figure 3.6. One tactor is placed on the fingertip of the thumb and another on the index finger to provide real-time vibrotactile feedback of the hand's position in space. This system can be used to either encode a positional error relative to a target in the workspace, or even a global positional signal to provide

information to the subject of where their hand is in space, hence the 'proprioceptive emulator' portion of the VIPER system name.

A current prototype of the device also includes a closed-loop feedback mechanism to accurately measure the timing and frequency output of the tactors in real-time. A solid-state 3-axis accelerometer is attached directly onto the motor casing, and data from the accelerometer is recorded at 1 kHz by the RT controller. The accelerometer is sensitive enough to not only precisely detect the onset and offset of the motor, but by applying spectral analysis techniques it can measure the precise frequency of operation at each point in time. Movement artifacts can easily be filtered out to isolate the vibratory signal. This technique was important for synchronization and alignment purposes, which will be discussed in a later section.

Figure 3.6: Vibrotactile feedback system (VIPER). Tactors are positioned on the fingertips of the thumb and index finger (green arrows) to apply continuous, graded vibrotactile feedback. Also shown is an active marker (red LED) attached to the back of the index finger, which is used by the motion capture system to track hand position.

3.1.3.3 General Purpose Output

Aside from the 6 analog channels currently being used for the VIPER system and other digital lines used for controlling graphical events, there are an additional 26 analog and 16 digital output lines available for other applications. One current use for one of the digital lines is to trigger a juice reward system for non-human primate experiments. Other possible applications of these lines include driving a larger number of tactors for more channels of vibrotactile feedback, generating or triggering auditory tones, or generating waveforms to provide electrical microstimulation for neurophysiological studies.

3.1.4   Synchronization & Data Storage

One of the most critical aspects of developing a real-time data acquisition and control system is ensuring that the system maintains deterministic control. Furthermore, when attempting to integrate multiple third-party hardware and software systems, the ability to properly align and account for any potential system delay is paramount. In this system, all possible latencies are meticulously calculated and stored along with the rest of the data. Although the RT controller operates deterministically at a set clock rate of 1 kHz, data from the motion capture system and the renderings of graphical output, for example, are not deterministic in nature. To account for these nondeterministic signals, additional synchronization techniques are used to accurately measure each of the system latencies.

The motion capture data has to initially be piped to the Host before making its way to the RT controller. This limitation was due to incompatibility of the system's proprietary

drivers/software with the real-time operating system. In practice, this introduces a slight delay in transmitting kinematics data to the RT controller to be processed, however, this latency is visually imperceptible to the subject. To circumvent any asynchrony in data storage, the motion capture system sends a hardwired 24-bit digital word directly to the RT controller containing the current sample frame number. This signal then allows the system to align the frame number received over the Ethernet connection to the RT controller's deterministic clock by shifting the Ethernet data to match the digitally transmitted frame number.

To achieve minimal graphics latency, a dedicated graphics computer is used to generate visual stimuli in the VR environment. Even with the dedicated GPU, there still exists a latency between the time the GPU receives a command to generate a graphical output and the time that graphical event takes place on the screens for a number of reasons. First, the GPU runs in a windows environment to take advantage of the powerful OpenGL graphics package, which will inevitably incur some latency and jitter. Second, there exists some latency in the graphics card that drives both monitors. Third, the LCD monitor refresh rate is only 60 Hz, so depending on when data was sent to the monitor, the actual image could be rendered at a variable point in time within the screen refresh window. To accurately measure the total graphics latency, a photodiode is placed in the corner of one of the LCD monitors out of the subject's field of view. Every state transition during each trial, a small visual stimulus at the location of the photodiode is toggled on or off. The voltage change induced in the photodiode by the stimulus is then fed back into the RT controller. Although the precise timing of the graphical display is not deterministic, by using the photodiode synchronization the actual time of visual stimulus presentation can be calculated and recorded for either online task flow control or offline alignment and analysis.

As described in Section 3.1.3.2, the vibrotactile feedback system (VIPER) is capable of providing closed-loop measurements as well. This is needed primarily due to the sluggish nature of the tactors. The RT controller is capable of sending a voltage command signal instantaneously, but there is a variable amount of spin-up time introduced by having to overcome the inertial load of the motor. By recording the accelerometer data attached to the tactor much like the photodiode for the graphics display, precise timing information about the tactor operation is known in real-time and is stored for offline analysis. Moreover, the accelerometer can also help calibrate each of the tactors by providing the precise frequency output for a given input voltage. The time-frequency analysis of the accelerometer attached to the tactor can be further analyzed offline to measure the precise motor frequency at each point in time.

Another possible source of data misalignment comes from the neural acquisition system. In order to align all of the raw neural data stored on the DSP controller, a synchronized digital time stamp is sent from the RT controller into the RZ-2 DSP and merged with the neural data stream. Even though the neural data is reliably sent directly to the RT controller via Ethernet (UDP) transmission, this synchronization allows for proper offline alignment of the raw waveforms stored on the DSP controller.

Finally, all signals that are received directly by the RT controller are by definition aligned properly to the 1 kHz clock. If this deterministic control of the RT controller is violated, an error code is presented to the researcher. Similarly, if any other system error occurs, including inter-system communication errors or hardware/software errors, an error code and description is presented on the Host GUI to alert the researcher.

Data are stored on the Host on a trial-by-trial basis. All of the recorded behavioral (motion capture, eye tracker) and neural data, trial parameters, state table information, and

synchronization information for a single trial are stored initially as a zipped collection of binary LabVIEW files. Raw waveforms from the neural data are stored locally on the DSP controller, but sorted spike detections are stored on the Host as well. A custom LabVIEW 'data translator' converts the binary LabVIEW files into a convenient Matlab structure format. The data translator can be run as experiments are in progress, allowing for real-time data analysis in Matlab on the Host computer.

3.2   SYSTEM EVALUATION

3.2.1   Capabilities

In its current state, the system is capable of obtaining data from multiple data acquisition sub-systems (neural data, limb position, eye position), deterministically process these data within the LabVIEW Real-Time environment, and generate graphical, auditory and vibrotactile stimuli output. Aside from third party hardware, the system is composed of mainstream, affordable, off-the-shelf PCs and National Instruments hardware. The custom LabVIEW software running within the Host/RT controller configuration provides a powerful yet adaptable platform to develop and implement experimental paradigms for a variety of potential research applications.

The Host software provides a straightforward user interface (Figure 3.3) for the researcher, with intuitively structured panels for task design and modification. Real-time displays of limb position, eye position, stimuli parameters, and other trial information are all available with the Host GUI. The flexibility of the state-machine architecture and simplicity of

the user interface allows for non-programmers to quickly implement novel experimental paradigms without writing any code whatsoever. Furthermore, changes to existing paradigms can even be made online during experiments to be implemented on the subsequent trial after the change is made, alleviating the need to stop an experiment to rewrite and recompile scripts as is required with other systems using conventional text-based programming languages.

The use of a Wheatstone stereoscopic display along with custom OpenGL and LabVIEW graphics software creates an immersive 3D virtual reality environment for the subject. The mainstream use of the OpenGL graphics package (e.g. video game development) allows for practically any virtual object to be rendered in the VR scene. The use of high-resolution graphics in combination with programmatic control of depth cues delivers a truly convincing perception of a 3D virtual environment (see Figure 3.5 for a single monitor screen capture example).

3.2.2 Performance

The system was designed to operate with minimal closed-loop latency (i.e. stimuli are generated as quickly as possible in response to a detected behavioral or neural event). A summary of the temporal performance characteristics of the system is shown below in Table 3.1.

Table 3.1: System latencies. Input latencies refer to the time between when the recorded event took places and went it was received by the RT Controller. Output latencies represent the time from RT Controller command to the actual output being generated. * ~ 18 msec is the latency of the motors when changing from a low to moderate frequency (65 to 90 Hz); the latency from completely off to a moderate frequency (0 to 90 Hz) is ~ 37 msec.

		Latency (milliseconds)
Input		
	Neural Data	< 2
	Motion Capture	< 16
	Eye Tracker	< 4
Output		
	Graphics	38 ± 8
	VIPER	18 ± 3 (37 ± 6) *

The system's ability to acquire 96 channels of sorted neural data reliably every millisecond was verified using simulated waveforms with known characteristics. The motion capture latency shown in Table 3.1 is a combination of the system hardware and software overhead in identifying and transmitting marker position data in addition to the slight latency introduced by passing through the Host. The accuracy of the motion capture system approaches sub-millimeter resolution depending on the volume of the workspace using a 6-camera setup as is currently implemented. During a system test, the 3D eye position accuracy was calculated to be less than 1 cm in the frontoparallel plane (x and y axes) and under 3 cm in depth (z axis).

Once all of the data has reached the RT controller, it is processed and output commands are generated within one clock cycle, which is strictly controlled by the RT controller's operating system. Using the photodiode feedback system described earlier, the graphics latency was characterized to be slightly more than two screen refreshes (~ 38 milliseconds) of the 60 Hz LCD monitors. This latency is imperceptible to the subject but nonetheless it is precisely recorded at each state transition for offline alignment. The latency of the VIPER system is variable depending on the state of the tactors: the latency from completely off to a high level of vibratory

stimulation takes a greater amount of time due to the initial burden of overcoming the inertia of the motors. Despite this latency, the VIPER system is still capable of providing near real-time, continuous vibrotactile feedback.

3.2.3   Current Implementation

Currently, the system is being used in non-human primate experiments and human psychophysical studies. The human studies have provided an excellent opportunity to verify system performance and gain user feedback. The system has demonstrated the ability to effectively track arm position with multiple markers, track 3D eye position, and generate a rich 3D VR graphics display. The VIPER system has been used to deliver real-time tactile stimuli as part of an ongoing study investigating the efficacy of vibrotactile feedback in sensory-motor control (see Chapter 4).

Multiple lab members, ranging from experienced programmers to those with non-technical backgrounds, have designed and implemented experimental paradigms from scratch without having an in depth knowledge of the underlying code that makes up the software. All users thus far have praised the system's ease-of-use and the ability to dynamically change experimental parameters. Overall, the system has successfully demonstrated the ability to perform highly sophisticated data acquisition with relatively seamless integration with third party hardware. Furthermore, the custom LabVIEW software has provided deterministic control using a mainstream platform with a large user base for support. Lastly, the VR and tactile displays have provided an immersive environment for subjects during experimentation.

4.0   HUMAN PSYCHOPHYSICS

4.1   INTRODUCTION

As the fields of brain-machine interfaces and human-computer interaction continue to advance (see Chapter 2), there continues to be a convergence in technology between neurophysiological recording methodologies and advanced computing interfaces. More specifically, the use of interactive devices to provide non-visual feedback for human-computer interaction has motivated researchers in the field of brain-machine interfaces to integrate such devices to provide additional sensory feedback for neural prosthesis control.

As described in Chapter 2, vibrotactile feedback has the potential be a simple yet effective method of providing somatosensory feedback for neural prosthetic applications. This approach has recently been employed for myoelectric-based prosthesis control in amputees where residual nerves from the amputated limb are implanted into a remaining functional muscle (Kuiken et al. 2009). Researchers found that not only could the user attain sufficient control over the prosthesis using the myoelectic signals generated from controlling the reinnervated muscle, but that users were also able to perceive tactile stimuli at the sight of nerve reinnervation. Based on this finding, researchers use this available sensory channel to deliver vibrotactile feedback to the user, providing tactile information from the prosthetic hand, which contained force sensors to

31

mimic tactile perception (Marasco et al. 2009). The end result was a prosthetic arm that utilized vibrotactile feedback as a sensory substitution method to provide the sense of touch to the user.

Inspired by the aforementioned study along with the years of work done on tactile perception, sensory-motor control and brain-machine interfaces (see Chapter 2), it was desirable to evaluate the potential application of vibrotactile feedback for neural prosthesis control. To this end, a two-part psychophysical study was conducted with able-bodied human subjects to investigate the efficacy of vibrotactile stimulation as a feedback control signal to guide reaching behavior in a virtual-reality environment. Moreover, this work intended to attain a further scientific understanding of multimodal sensory integration and motor control.

4.2   EXPERIMENT 1: VIBROTACTILE MOTOR CONTROL

In the first of two experiments, the aim was to evaluate the utility of a simple two-tactor vibrotactile display (described in Chapter 3) as a real-time reach kinematics feedback controller. To do so, motor performance was measured as subjects performed a reach task in a virtual-reality environment under different visual and tactile feedback conditions. By comparing motor performance with and without various visual cues as vibrotactile feedback was toggled on or off, the goal was to determine the relative effectiveness of the vibrotactile signal to provide useful sensory information in controlling arm movements.

4.2.1 Methods

4.2.1.1 Experimental Setup

Nine able-bodied, sighted human subjects performed a trajectory-following reach task in a virtual reality environment while wearing a glove that provided vibrotactile feedback. The experimental setup is described in greater detail in Chapter 3. For this study, subjects sat in front of a Wheatstone stereoscopic display. Subjects were not able to physically see their hand using this setup, but instead hand position was rendered as a sphere in the VR environment. Although the VR display is capable of providing realistic 3D graphics, for this experiment all visual stimuli (including rendered hand position) were confined to a 2D plane in the 3D workspace, essentially creating a 2D VR display. Behavioral control, data acquisition and graphics display were all performed using custom LabVIEW software and third party hardware (see Chapter 3). Subjects wore a glove on either their reaching (dominant) or non-reaching (non-dominant) hand. An active marker (LED) was placed on the back of the glove to track hand position in the 3D workspace. Vibrotactile actuators were attached to the thumb and index finger as described in Chapter 3.

4.2.1.2 Task Conditions

Subjects performed a trajectory-following reach task under different visual and tactile feedback conditions. Every trial began with the presentation of a start target located near the bottom of the VR workspace. After the subject hit the target (i.e. moved the hand cursor onto the target) an instructed trajectory and end target were displayed. The trajectory is displayed as a tube connecting the start and end target in the VR environment. All task parameters were randomly

33

interleaved during the experiment. There were 7 possible end targets, all in an upward or oblique direction relative to the start target. Three possible trajectories could be displayed (straight, curved left or curved right). See Figure 4.1 for an illustration of start target and end target locations, and example instructed trajectories. All targets were placed along semi-circular arc 17.5 cm from the start target. Subjects were instructed to reach from the start target to the end target along the specified trajectory as accurately as possible. They were given 8 seconds to complete each trial. If the subject's hand deviated too far from the trajectory (3 cm) or took too long to complete the reach (> 8 seconds) they failed the trial. Failed trials were not considered in the analysis. Vibrotactile feedback provided a real-time, directional error signal relative to the instructed trajectory in a proportional fashion: if the subject's hand position was to the left, or counter clockwise, relative to the instructed trajectory the thumb vibrated to signal the subject to move to the right to get back along the correct path. Conversely, vibration was applied to the index finger if the subject's hand position was to the right, or clockwise, relative to the trajectory. The further the subject deviated from the correct path, the more intense the vibratory stimulus became. Consequently, the subject's goal was to minimize the vibratory stimulus by following the instructed trajectory. To ensure continuous vibrotactile feedback, subjects had to essentially be perfect in order to receive no tactile feedback: the 'dead zone' or error region where no vibration was delivered was only 1 millimeter surrounding the instructed trajectory. Thus, the subject almost always received at least a minimal level of vibration encoding their current directional reach error.
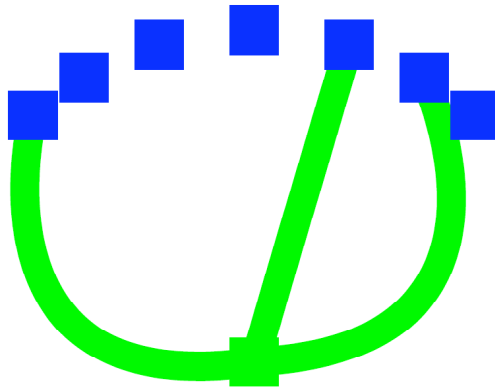
Figure 4.1: Illustration of example target and trajectory configurations. The start target (green square) always remained at the same location. There were 7 possible end targets (blue squares) with 3 possible trajectories (green tubes) to each of the end targets. Three example trajectories are shown, one for each type of trajectory (straight, curved left, curved right).

In addition to the variable target and trajectory presentations, three forms of feedback were randomly toggled on and off during the experiment to create a total of 8 experimental conditions. Visual feedback of the hand cursor, visual feedback of the instructed trajectory and end target, and vibrotactile feedback of hand positional error were either displayed or not for a given trial. Table 4.1 below shows a logic table and list of abbreviations for each of the experimental conditions.

Table 4.1: Summary of task conditions. 1 = feedback is given, 0 = not given. V = vibrotactile (also abbreviated as Vib.), T = trajectory/end target, C = cursor, N = none.

Abbreviation	Vibrotactile	Trajectory & End Target	Cursor
VTC	1	1	1
TC	0	1	1
VC	1	0	1
C	0	0	1
VT	1	1	0
T	0	1	0
V	1	0	0
N	0	0	0

Each form of feedback was randomly selected for each trial within an experimental session but never changed during a single trial. For all trial types, visual feedback of the cursor and start target was always given until the subject hit the start target. For trials when the trajectory and end target were not visible, an initial but brief presentation of the trajectory and end target was given (150 milliseconds) immediately after the start target was reached. For trials where the cursor was not visible, it also disappeared just after start target acquisition and before the subject began the reach along the instructed trajectory. Vibrotactile stimulation was engaged immediately before the reach began for the trials where vibrotactile feedback was delivered. Figure 4.2 below shows two examples of visual feedback for hand cursor and trajectory visible task conditions.
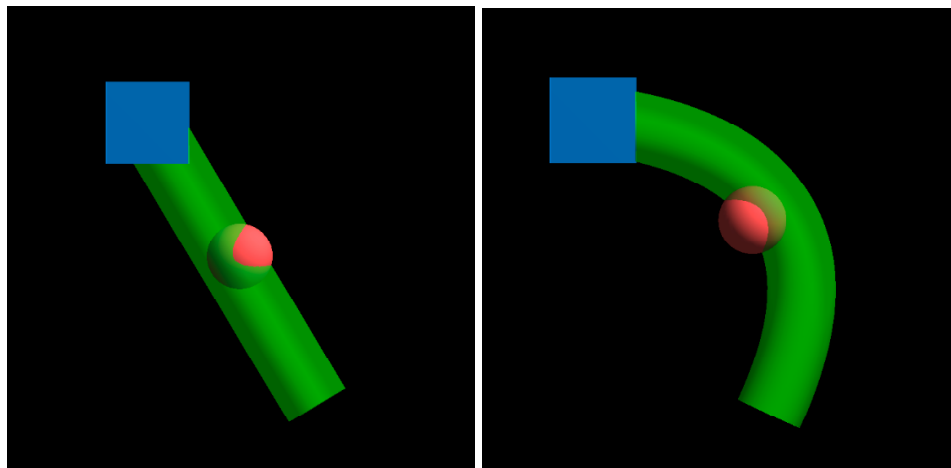

Figure 4.2: Straight and curved trajectory examples for the full vision condition. The red sphere represents the subject's hand position as they reach along the instructed trajectory (green tube) toward the blue end target.

Note that the hand cursor is the same diameter as the tube, which provides the subject with high visual acuity to detect reach errors during the experiment under visual feedback conditions. Occlusion cues provide the perception of 'entering' and 'exiting' the tube as the subject reaches along the trajectory.

4.2.1.3 Performance Metrics

To evaluate subject performance, multiple reach error metrics were used. Subjects were instructed to be as accurate as possible, without necessarily an emphasis on making fast reaches. The following three metrics were calculated to measure subject's reach accuracy:

1) Movement Error (reach root-mean-square error)

$$ME = \frac{\sum |e_i|}{n} \quad \text{where } e_i = \sqrt{\left(error_{x,i}\right)^2 + \left(error_{y,i}\right)^2}$$

e_i represents the magnitude of the point by point error defined by the perpendicular distant to the instructed trajectory at the ith time point.

2) Movement Offset (reach bias)

$$MO = \bar{e}$$

3) Movement Variability (reach variance)

$$MV = \sqrt{\frac{\sum (e_i - \bar{e})^2}{n - 1}}$$

These particular metrics were adapted from MacKenzie et al. (2001) to provide insight into not only a classical error metric (root-mean-square error), but also the more specific nature of the error by decomposing the error into its constitutive bias and variance (i.e. offset and variability) components. Reach duration, defined as the time to complete the reach, was also used to evaluate performance.

4.2.2   Results

Figure 4.3 shows the performance results for the nine subjects who participated in the study using the three error metrics described above (ME, MO, MV). Subjects performed on average a total of 211 trials, resulting in approximately 26 trials for each of the 8 feedback conditions.
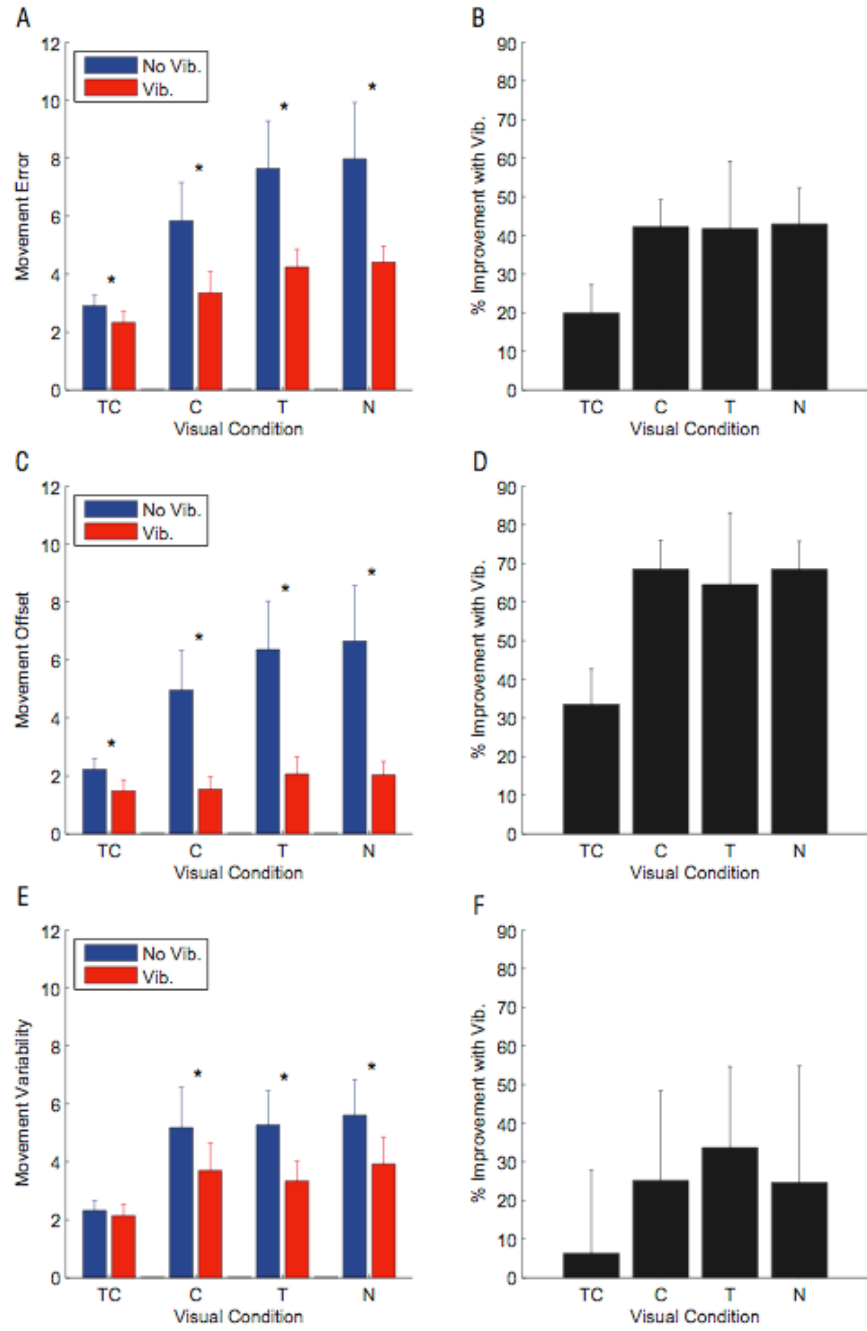
Figure 4.3: Summary of subject performance for the 2D trajectory-following reach task. Subject error using Movement Error (A,B), Movement Offset (C,D), and Movement Variability (E,F) metrics. For A,C and E, Error is shown for each visual feedback condition: TC = trajectory and cursor visible, C = only cursor visible, T = only trajectory visible, N = nothing visible. No Vib. = without vibrotactile feedback (blue), Vib. = with vibrotactile feedback (red). Panels B,D and F show the relative percent improvement in performance (i.e. reduction in error) for the data shown in A,C, and E. Data points represent population means (n=9); error bars represent inter-subject standard deviations. * All pair wise comparisons were statistically significant (paired t-test; $p < 0.05$) except for movement variability in the TC condition.

As shown in Figure 4.3, the overall effect of the vibrotactile feedback was to reduce reach error, thereby improving motor performance. The data in Figure 4.3 is also given in Table 4.2, along with p-values for each of the Vib.-NoVib. pair wise comparisons. Movement error (panels A,B) was significantly reduced for all visual conditions under vibrotactile feedback guidance. This reduction was greater when at least one form of visual feedback was missing (C,T, and N), as is illustrated by the increase in percent improvement in Figure 4.3 B. Also note that in both the Vib. and No Vib. case, errors increased as visual information was removed (i.e. going from having both forms of visual feedback, to one, to none). This implies that both visual and tactile information are integrated (when given) to guide movements in this task. Moreover, the fact that performance in the full vision case (TC) was improved with vibrotactile feedback further suggests a blending of sensory modalities as opposed to choosing one or the other.

Figure 4.3 C and D show the results for movement offset (i.e. movement bias) while panels E and F show results for movement variability (i.e. movement variance). Movement offset was relatively constant regardless of the visual feedback condition under vibrotactile guidance, implying that the tactile feedback was sufficient in reducing movement offset without visual feedback. Movement variability was also reduced in the vibrotactile feedback conditions, but not nearly to the extent that movement offset was reduced. This suggests that movement error, which is composed of movement offset and movement variability, is primarily reduced by vibrotactile feedback due to a decrease in movement offset. A reduction in movement offset shows that subjects could effectively find the underlying trajectory while a reduction in movement variability demonstrates that the subjects reduce their oscillation around the underlying trajectory. Note that all but one pair wise comparisons shown in Figure 4.3 show that

vibrotactile feedback significantly reduced each type of error. The only non-significant comparison was in the full vision case with movement variability (see Table 4.2).

Table 4.2: Statistical summary of the 2D trajectory-following reach task. Data presented as population mean $\pm$ standard deviation. P-values (paired t-tests) are given for each Vib.-No Vib. comparison for each visual feedback condition. (n = 9).

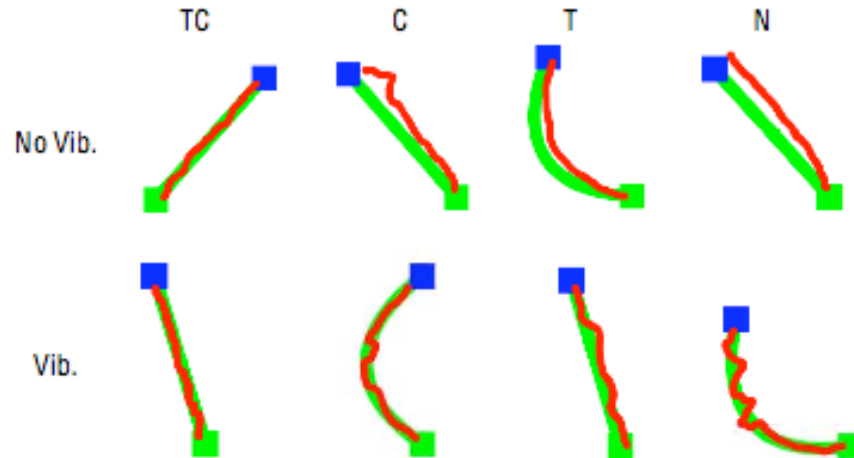|		Visual Condition			
		TC	C	T	N
Movement Error	No Vib.	2.9 ± 0.4	5.8 ± 1.3	7.6 ± 1.6	8.0 ± 1.9
	Vib.	2.3 ± 0.4	3.4 ± 0.7	4.2 ± 0.6	4.4 ± 0.6
	p	**0.0024**	**7.1e-5**	**1.2e-5**	**3.6e-5**
Movement Offset	No Vib.	2.2 ± 0.4	5.0 ± 1.4	6.4 ± 1.7	6.7 ± 1.9
	Vib.	1.5 ± 0.4	1.5 ± 0.4	2.1 ± 0.6	2.0 ± 0.5
	p	**3.6e-4**	**1.2e-6**	**1.0e-6**	**1.3e-6**
Movement Variability	No Vib.	2.3 ± 0.3	5.2 ± 1.4	5.3 ± 1.2	5.6 ± 1.2
	Vib.	2.1 ± 0.4	3.7 ± 1.0	3.3 ± 0.7	3.9 ± 0.9
	p	0.1615	**0.0092**	**0.0003**	**0.0025**

Figure 4.4: Example trials for each of the 8 experimental conditions. Red traces represent the path of the reach from the start target (green square) along the instructed trajectory (green tube) to the end target (blue square). Top row: without vibrotactile feedback. Bottom row: with vibrotactile feedback. From left to right: TC, C, T, N.

Figure 4.4 shows an illustration of single representative trials for each of the 8 task conditions. In the full vision case, subjects made very accurate movements. The improvement in performance with vibrotactile feedback originates from the slightly improved ability to precisely follow the exact underlying trajectory. In the cursor only conditions, subjects were able to realize they need to make a corrective movement using visual feedback of the cursor and their motor plan, while the addition of vibrotactile feedback kept the subject accurately following along the correct trajectory, but did not eliminate small deviations from the trajectory. In the trajectory vision only conditions, a view of the trajectory along with vibrotactile feedback was again sufficient to maintain an accurate path, however, without the tactile feedback subjects tended to deviate from the path but would eventually hit the end target. Lastly, under the null vision case, subjects would try to rely on an open-loop motor plan but often strayed off course. The vibrotactile feedback was able to allow the subjects to identify the underlying trajectory, but due to the lack of visual feedback subjects tended to oscillate around the correct trajectory as the

error signal pushed them back and force along the path. It is also important to note that a large amount of variability existed in the non-vibrotactile feedback trials when subjects had to wave around near the end target (not shown in Figure 4.4) in order to hit it. This typically never happened under vibrotactile guidance.

Due to the fact that the error metrics used to evaluate performance are measurements that average across a given reach, reach duration (i.e. time to complete the reach) was also calculated for each task condition as shown in Figure 4.5.
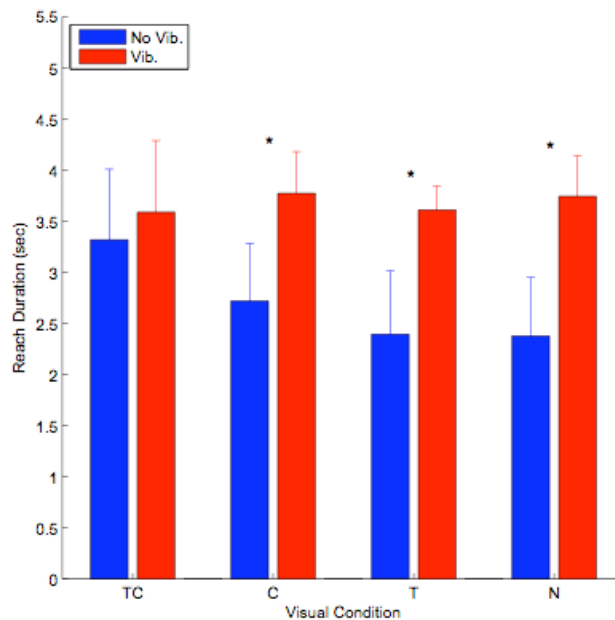
Figure 4.5: Reach Duration for each of the task conditions. Reach Duration is defined as the time to complete the reach (i.e. the time from the acquisition of start target to the time of end target acquisition).

As Figure 4.5 shows, subjects tended to reach slower the more feedback they received, with vibrotactile feedback acting as the rate-limiting factor. Regardless of visual feedback, subjects tended to make the same speed reach while under vibrotactile feedback guidance. However, as less and less feedback was given (i.e. moving from left to right in the No Vib.

conditions) subjects tended to increase the speed of their reaches. Subjects were instructed to be as accurate as possible, and there was no requirement on speed as long as they met the loose time constraint of 8 seconds. With that in mind, it appears that the strategy that most subjects took was to reach as slowly as was required in order to effectively use whatever type or amount of feedback that was provided.

One concern that could be raised by fact that subjects made slower reaches with vibrotactile feedback is that the enhancement in performance gained by using the vibrotactile feedback is a consequence of forcing subjects to slow down and pay more attention to their reach, rather than providing an additional, useful feedback control signal. To alleviate this worry, a correlation analysis was performed within each condition for each subject to identify any relationship between reach error and reach duration. The reasoning behind this approach is as follows: if going slower resulted in smaller movement error, then this relationship should show up within each of the specific conditions for each subject. Figure 4.6 shows the Pearson's correlation coefficient between reach duration and movement error within each task condition for each of the 9 subjects.
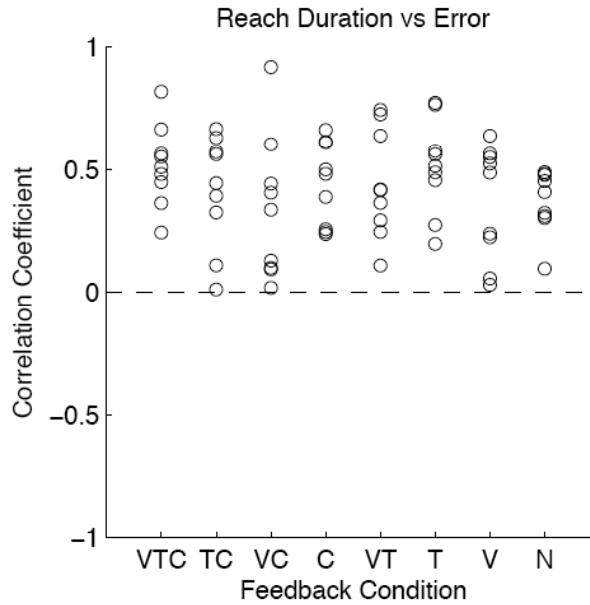
Figure 4.6: Correlation between reach duration and movement error. Correlation between Reach Duration and Movement Error is given for within each task condition for each subject. Each data point represents the Pearson's correlation coefficient of each subject for a given feedback condition. Note that none of the correlation coefficients are negative, implying that there is no improvement in task performance for slower reaches. In fact, subjects tended to do worse when making slower reaches. Feedback Conditions: V = vibrotactile, C = cursor, T = trajectory/end target, N = none.

As shown in Figure 4.6, subjects actually tended to make larger errors when reaching slower (i.e. longer reach duration), as indicated by the positive correlation values. This shows that the better performance using vibrotactile feedback is not just due to the longer reach duration, but rather the additional information the feedback signal provides. See Figure A1 in Appendix A for individual linear fits for each subject within each condition used to characterize the relationship between reach duration and error described in Figure 4.6.

Out of the 9 subjects, 4 of the subjects received vibrotactile feedback on their non-reaching hand while 5 subjects received vibrotactile feedback on their reaching hand. Figure 4.7 shows a comparison between the non-reaching (off hand) and reaching-hand groups.
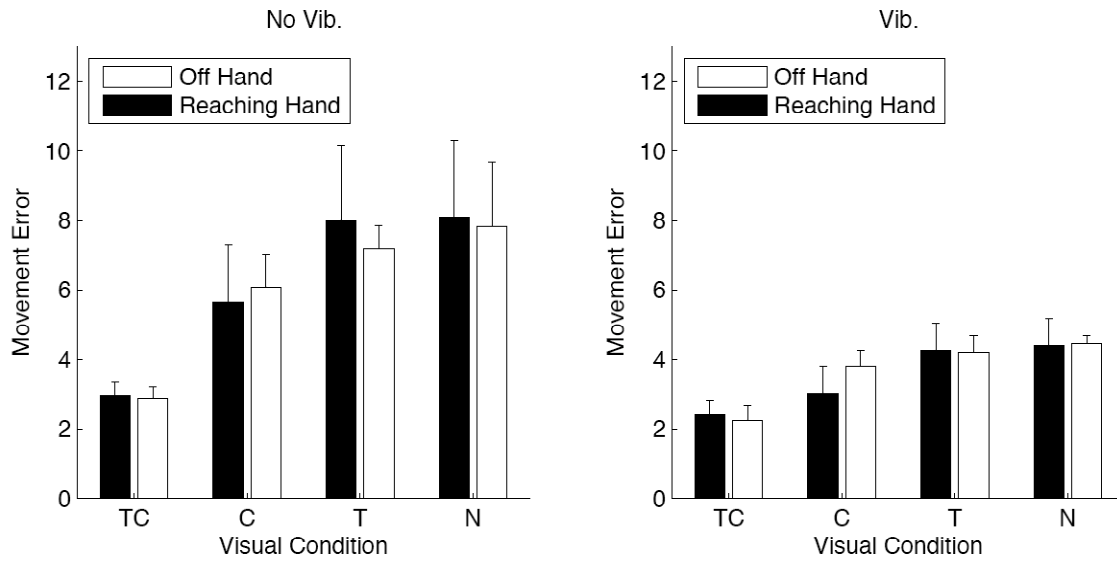
Figure 4.7: Comparison of movement error between off hand and reaching hand groups. Comparison is shown across visual conditions without vibrotactile feedback (left) and with vibrotactile feedback (right). There was no statistical difference between the two groups.

There were no statistical differences in any of the visual feedback conditions for either the no vibrotactile or vibrotactile case (all p-values > 0.05 using paired t-tests, n = 4 and n = 5 for the off hand and reaching hand groups respectively). Assuming equality between the two groups, all of the data preceding Figure 4.7 has pooled the two groups together.

4.2.3 Discussion

The goal of this first experiment was to be able to characterize the ability of subjects to use a simple vibrotactile display to dynamically guide reaches in a VR trajectory-following task under different feedback conditions. Overall, subjects were able to utilize the vibrotactile signal to partially recover the drop in performance (i.e. increase in error) due to missing visual feedback, and were even able to enhance natural performance in the full vision case (see Figure 4.3 A and

B). The trend of increasing error as visual feedback is removed, especially in the vibrotactile feedback conditions, suggests that subjects integrate both visual and vibrotactile modalities to guide their arm movements. The nature of the vibrotactile feedback signal is a positional error signal, which could be thought of as a sensory substitution signal for an enhanced form of proprioception. It is important to note that proprioceptive feedback was present in all task conditions. The fact that subjects show decrease in performance without visual or vibrotactile feedback implies that both the visual and vibrotactile modalities provide reliable information to be used to guide closed-loop movements for this task.

The appeal of this type of task as compared to a standard point-to-point reaching task was that during the reach, the path that subjects aim to take is well defined. This allows for precise calculation of movement error at each point in time along the duration of the reach. Moreover, by instructing the exact trajectory subjects must follow, each reach is performed in a tightly constrained fashion, allowing for even subtle changes in error patterns to be detected across conditions. Also, the experimental design of randomly interleaving any type of target/trajectory or task condition alleviates any concern of learning effects or inherent bias in task-specific performance. The goal of designing the paradigm in this fashion was to prevent any memory-guided reaches, forcing the subjects to rely on the online feedback being presented for a given trial.

The use of vibrotactile displays has been studied extensively over the past few decades (see Chapter 2). However, there has yet to be a study that uses this type of sensory feedback to dynamically guide movements as a real-time control signal with the level of sophistication achieved here. Lieberman et al. (2007) used vibrotactile feedback to have subjects match arm joint angles with a virtual arm presented on a screen. The goal of their study was similar in that

the vibrotactile feedback served as a proprioceptive enhancement technique, but they did not use the sensory signal to provide continuous feedback to guide movements as in this experiment. Moreover, the ability of subjects to learn to effectively use this simple and intuitive form of sensory feedback demonstrates the potential of such devices to enhance sensory-motor control for a variety of applications, as will be discussed in Chapter 5.

4.3   EXPERIMENT 2: DEPTH PERCEPTION AND LEARNING

In the first experiment (section 4.2 above), subjects performed a trajectory-following reach task along trajectories confined to a 2D plane in the VR environment. This was done to keep the task simple to attain a clear understanding of the utility of vibrotactile feedback. Task conditions, including whether or not vibrotactile feedback was given, were randomly interleaved throughout an experimental session to avoid biasing task-specific performance with learning effects.

Studies have shown that visual acuity is more accurate than proprioception in the frontoparallel plane (azimuth), whereas proprioception actually provides greater acuity in depth than does vision (Scott & Loeb 1994, van Beers 2002). With this in mind, along with the notion that vibrotactile feedback could be used as a sensory substitution method for proprioception as in experiment 1, one of the primarily goals of this second experiment was to see to what extent vibrotactile feedback could enhance depth perception in a 3D version of the trajectory-following task used in experiment 1. The second goal of this experiment was to study motor skill learning

associated with using a vibrotactile display by using a block experimental design rather than randomly interleaving task conditions.

4.3.1   Methods

Similar to experiment 1, subjects performed a trajectory-following reach task where they were instructed to reach from a start target to an end target along an instructed trajectory as accurately as possible. The same performance constraints were applied: subjects had to complete each reach within a soft time constraint of 8 seconds and could not deviate from the instructed trajectory at any point in time more than 3 cm. For this experiment, there was always visual feedback of the hand cursor, trajectory and end target. There were 27 possible end target locations, forming an evenly sampled 3D sphere of possible targets surrounding the start target. There were also three different trajectory types: straight, single inflection (C shape) and double inflection (S shape) tubes. The single and double inflection tubes could also have one of four possible longitudinal rotations applied: 0, 90, 180, or 270 degrees along the tube's long axis. Screen captures of a few example trials are shown in Figure 4.8 below.

Vibrotactile feedback was delivered using the same two-tactor glove as in experiment 1, except for this experiment the tactors encoded error in depth, rather than in the x-y plane. The mapping of the error signal was as follows: if a subject moved too far in depth (away from their body or positive in the z axis) the tactor on the index finger would vibrate, indicating that their hand is too far and that they need to make a corrective movement toward their body. Conversely, if the subject's hand was too close to their body, the tactor on the thumb would vibrate to inform the subject to make a corrective movement away from their body. This

49

encoding scheme did not encode error in either the x or y axis (azimuth). Also, the tactors provided a continuous graded error signal as before that was proportional to the magnitude of their error such that as a subject moved away from the instructed trajectory, the frequency (intensity) of the vibration increased. All subjects wore the tactor glove on their reaching hand.

Two groups of subjects were used for this experiment: a test group (n = 4) and control group (n = 3). The test group performed a block of 50 trials without vibrotactile feedback, followed by a block of 100 trials with vibrotactile feedback, and then ending with another block of 50 trials without vibrotactile feedback. Control subjects performed 200 trials without vibrotactile feedback, but took short breaks at the same intervals as the test subjects. Movement error (as defined in section 4.2.1.3) was used to characterize task performance.
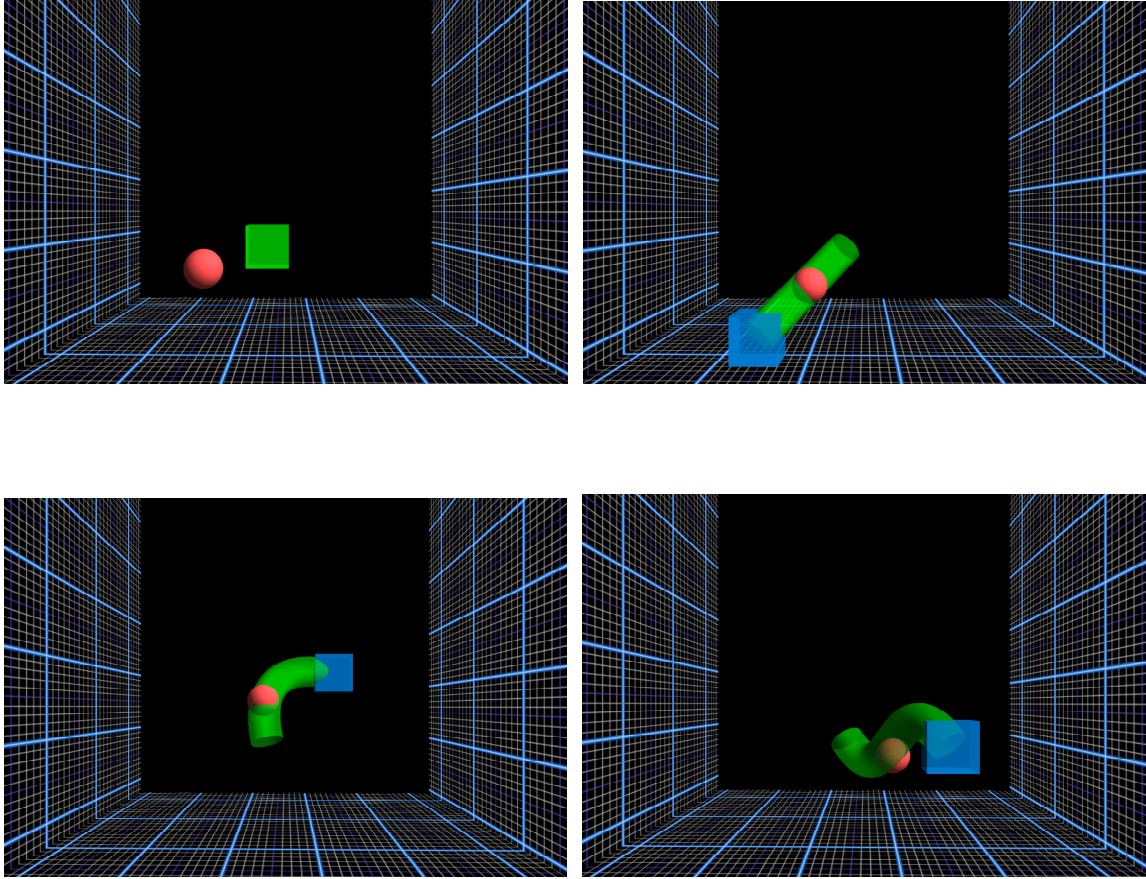
Figure 4.8: A few example screen captures of the 3D trajectory-following task. Top left shows the beginning of a trial, where subjects acquire the start target (green cube). The other panels show example reaches for different trajectory types (straight, single inflection and double inflection). Subjects reach along the green tube to the blue end target using the light red sphere as visual feedback of their hand position in the 3D workspace.
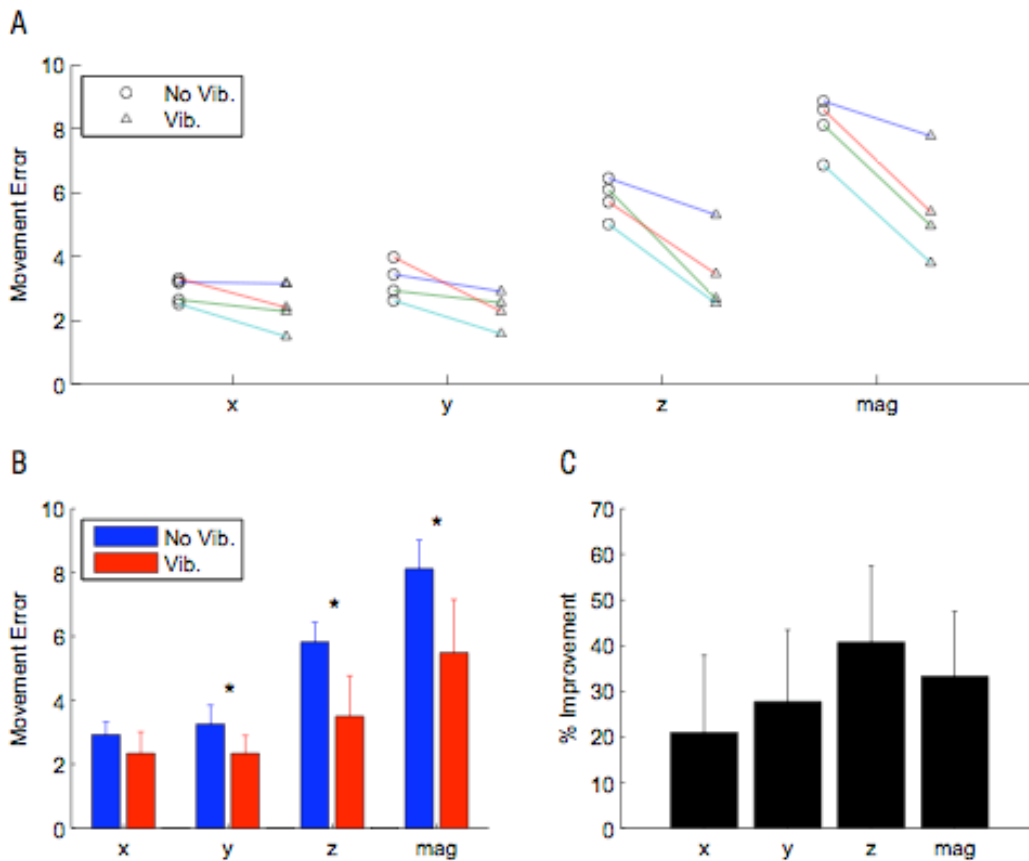
4.3.2 Results

Figure 4.9: Performance summary for test subjects in the 3D trajectory-following task. Movement error is shown individually by each axis (x,y,z) and overall magnitude. A: individual subject performance with and without vibrotactile feedback. Each subject is color coded (n = 4). Circles represent movement error without vibrotactile feedback. Triangles represent error with vibrotactile feedback. A negative slope indicates a reduction in error with vibrotactile feedback. B: Population results with and without vibrotactile feedback. C: Percent improvement (i.e. reduction in error with vibrotactile feedback) for each axis and overall magnitude (mag).

For this task, movement error was calculated using the same method as in experiment 1 (represented as magnitude or 'mag' in Figure 4.9). In addition, since vibrotactile feedback only encoded error in the z-axis, it was desirable to analyze the movement error for each axis independently. The overall effect of the vibrotactile feedback was to significantly decrease

52

movement error. As shown in Figure 4.9 A, each subject showed a reduction in the magnitude of error with vibrotactile feedback. In comparing the individual axis errors, subjects without vibrotactile feedback made larger errors in depth than in x or y as predicted by the assumption that visual acuity in depth is not as accurate as in azimuth. Furthermore, as shown in Figure 4.9 C, vibrotactile feedback showed the greatest benefit in the z-axis compared to x and y. Due to the inherent coupling of the coordinate axes, the depth error signal encoded by the tactors also provided a performance improvement in the x and y axes as well, although the improvement in the x-axis was not statistically significant. See table 4.3 below for a summary of the data shown in Figure 4.9.

Table 4.3: Statistical comparison of movement error for each axis and overall magnitude in the 3D task. P-values were obtained from paired t-tests between no vibrotactile and vibrotactile conditions (n = 4). There were significant reductions in error in overall magnitude, in addition to the y and z axes, but not the x axis.

	x	y	z	mag
No Vib.	2.9 ± 0.4	3.2 ± 0.6	5.8 ± 0.6	8.1 ± 0.9
Vib.	2.3 ± 0.7	2.3 ± 0.6	3.5 ± 1.3	5.5 ± 1.7
p	0.0920	**0.0335**	**0.0083**	**0.0160**

The benefit of the vibrotactile feedback in depth was great enough to reduce the z-axis error to a level comparable to the x and y errors for the no vibration case. The net effect of this drastic reduction in depth error along with the inherent coupling of the axes is to reduce the magnitude of movement error by over 30 percent (Figure 4.9 C).

The second aim of this experiment was to study potential learning effects associated with using vibrotactile feedback. Figure 4.10 below shows the block-by-block performance and relative improvements for both test and control groups.
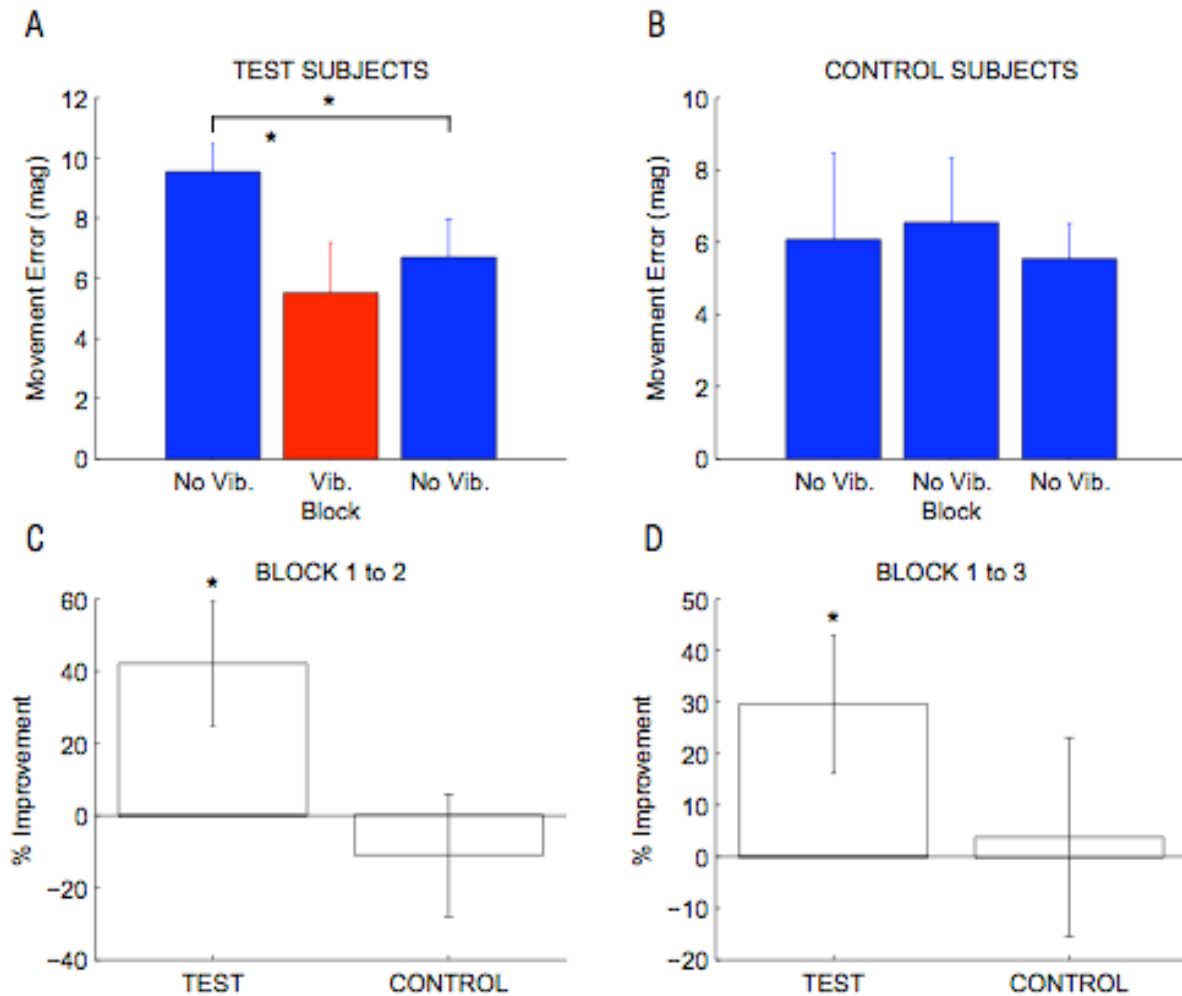
Figure 4.10: Block-by-block performance and learning effects for test and control groups. A: magnitude of movement error for the test subjects for each block (n = 4). B: magnitude of movement error for the control subjects (n = 3). C: Relative improvement from block 1 to block 2 for test and control subjects. D: Relative improvement from block 1 to block 3 for test and control subjects represented as a percent change.

As shown in Figure 4.10, test subjects not only showed a significant reduction in error while using the vibrotactile feedback, but they also showed a significant residual learning effect

in block 3. See Table 4.4 below for statistical comparisons. This is evident in the significantly

lower error in block 3 compared to block 1 for the test subjects (Figure 4.10 A) and is also shown

Figure 4.10 D as percent improvement. This residual learning effect was not present in control

subjects, as shown in Figure 4.10 D and in Table 4.4 ($p = 0.3691$ for control subjects while $p =$

0.0060 for test subjects in comparing error from block 1 to block 3). This suggests that the use

of the vibrotactile feedback helped subjects improve their general motor skill in reaching in this

3D virtual environment. It is important to note that any learning effects must generalize to the

motor skill of reaching in the VR workspace, rather than any target or trial specific learning.

This is due to the extremely low chance of a subject ever repeating the same exact trial enough

times to show learning effects. With 27 possible targets, 3 different tube types with 4 possible

different rotations, and two feedback conditions (vibrotactile or no vibrotactile feedback) there

exist 648 possible permutations of conditions for a 200 trial experimental session.

Table 4.4: A statistical summary of learning effects for control and test subjects. P-values show comparison of block learning from block 1 to block 3, calculated using paired t-tests. Test subjects showed significant learning ($p < 0.05$) while control subjects did not.

		Control	Test
	1	6.1 ± 2.4	9.5 ± 1.0
Block	2	6.5 ± 1.8	5.5 ± 1.7
	3	5.5 ± 1.0	6.7 ± 1.3
Block 1 to 3	p	0.3691	**0.0060**

4.3.3   Discussion

This second experiment aimed to build upon the first to address two concepts left unattended, namely: the ability of vibrotactile feedback to enhance performance in a 3D reach task, and to characterize any learning effects associated with using vibrotactile feedback. By encoding error only in depth using the two-tactor feedback device, subjects showed drastic reductions in depth error, proving that vibrotactile feedback can indeed provide a more than sufficient sensory signal to enhance performance.

The unexpected benefit of the vibrotactile feedback came in the reduction in error in the x and y axes as well. Although the improvement in the x axis was not significant, there seemed to be a trend, and the significant improvement in the y axis clearly shows that although the vibrotactile feedback signal only encoded depth error, the effect was to reduce error in all dimensions. This can be explained by the inherent coupling of the coordinate axes and by the perspective view of the subjects of the workspace. For example, if a subject were making an oblique reach toward the target located at the corner of a cube from the center and had relatively poor depth perception compared to azimuth, subjects would tend to make erroneous corrective movements in azimuth to eliminate the perceived error. So if a subject were too far in depth but did not necessarily notice this error, they would end up creating an error in azimuth to make the cursor line up with the trajectory from their perspective view of the 3D scene. Put another way, it makes it much easier to be accurate in the x-y plane when accurate in depth.

As in experiment 1, natural proprioception was also present on every trial, so again the enhancement in performance shows that vibrotactile feedback used in this context can provide sensory information capable of exceeding natural proprioceptive ability. Furthermore, since visual feedback was also always present, the vibrotactile feedback modality served to enhance

overall sensory-motor skill. It appears that subjects are able to incorporate this novel sensory modality with relative ease, and integrate it with visual and proprioceptive information to attain a higher level of motor performance.

To address the issue of learning, a block-to-block comparison seems to be sufficient to show the apparent residual learning effects after using vibrotactile feedback, implying that motor skill consolidation took place. As a result of the variable nature of each trial due to the novel target and trajectory presentations, trial-by-trial learning effects were not clear like in other psychophysical studies. Nonetheless, the data shown in Figure 4.10 suggest that subjects seem to consolidate a higher resolution internal model of the workspace aided by the additional sensory modality such that even after vibrotactile feedback has been taken away, they are able to perform at a higher level of accuracy. To study more subtle learning effects such as learning rates, more subjects would be needed and perhaps a modification in the paradigm would also be necessary. With only 4 and 3 subjects in the test and control groups respectively, it does cast some doubt on the power of the learning effects shown in Figure 4.10. Further experiments are underway to more formally address these issues.

5.0   CONCLUSIONS & FUTURE WORK

In an era where the line between different disciplines continues to blur, it is at the intersection of two cutting edge areas of research where the focus of this thesis lies. Technological developments in the field of brain-machine interfaces in conjunction with advances in modern computer interface systems have the potential to lay the foundation upon which sophisticated medical devices could one day provide viable therapies for currently unmet clinical needs. To this end, a general-purpose, state-of-the-art data acquisition, control and feedback system was developed for brain-machine interface and human psychophysics research. This system features seamless integration with third party data acquisition systems, deterministic system control, a highly powerful yet easily adaptable user interface, a rich 3D virtual reality display, and a novel vibrotactile feedback device.

One of the most important factors in developing an effective technological system such as the one presented here is the foresight to develop a platform capable of growing and adapting to needed upgrades and further system development. By building a system using off-the-shelf hardware and LabVIEW based software, any upgrades in technology or addition of functionality will be relatively easily integrated into the system. The large community of LabVIEW users provides a common language to share resources and get support. LabVIEW also provides an unmatched ability to seamlessly integrate with third party hardware and software systems. It is often the case, especially in the research community, that technology developed for a particular

application is the technology that is used for the next decade or more. In the development of this system it was one of the top priorities to create a modular, adaptable system capable of growth and future development without requiring a complete system overhaul.

As described in Chapter 2, one of the crucial areas for development in brain-machine interface research is to provide useful somatosensory feedback to emulate a more natural form of motor control of a neural prosthesis. The psychophysical experiments presented in Chapter 4 aimed to provide a proof of concept to motivate the use of vibrotactile feedback devices to supply this desired sensory modality for brain-machine interface applications. The results of these studies showed that vibrotactile feedback could replace or even enhance visual feedback in a virtual-reality reaching task. Moreover, when used in a 3D virtual reality environment, even a simple, two-channel continuous vibrotactile display has the ability to greatly enhance motor skill and motor learning as well.

These experiments have also provided basic scientific insight into sensory-motor control. Subjects were able to quickly learn to effectively integrate the vibrotactile feedback as a useful sensory channel for dynamic arm movement control. One subject even described the experience of using this device as an invisible force field that guided the hand along the trajectory during the task. In essence, subjects appeared to embody this new sensory modality and rely on its information content in an integral fashion, much like how we naturally blend our visual and proprioceptive senses in controlling our arm movements. It is this concept of embodiment where the true potential of this type of feedback lies: if a user of a neural prosthesis were to receive vibrotactile feedback of the prosthesis kinematics, the hypothesis would be that with training the user would not necessarily think of the stimulus as a vibration, but rather he or she would

embody the sensory modality and interpret it as the true 'proprioceptive' feedback that it represents.

Even with a potential proof of concept in using vibrotactile feedback for brain-machine interface applications, the question still remains as to how users will integrate such feedback. Furthermore, depending on the particular disease or injury of the patient, vibrotactile feedback may not be a viable sensory substitution method for providing somatosensory feedback due to paralysis or lack of tactile perception. Work by Romo and colleagues (Romo et al. 2000) has shown that direct cortical microstimulation in the somatosensory cortex of non-human primates is indistinguishably perceived as a vibratory stimulus on the skin. Thus, a similar type of somatosensory feedback could be given to the subject via direct electrical stimulation of somatosensory cortex rather than a vibrotactile display on the skin. Numerous researchers have investigated stimulating the peripheral nervous system to provide sensory feedback for prosthesis control (e.g. Dhillon & Horch 2005). Recently researchers have even begun microstimulating in somatosensory cortex in non-human primates to evoke some behavior in response to the stimulus (London et al. 2008). This work is currently very rudimentary, but results so far have shown that non-human primates are able to effectively discriminate various stimulation patterns delivered to somatosensory cortex through microstimulation, showing promise for its potential use as a somatosensory feedback method for neural prosthesis control.

With the work presented in this thesis, the hope is that future users of this system will be poised to develop high-performance, closed-loop neural prostheses complete with immersive visual feedback along with somatosensory feedback given through vibrotactile stimulation and eventually cortical microstimulation. Currently, microstimulation capability is being developed for the system for multi-channel stimulation applications. Furthermore, with a system capable of

recording limb movements and eye movements while generating rich VR graphics and tactile feedback, a wide variety of psychophysical studies could be conducting investigating concepts such as eye-hand coordination, multimodal sensory integration, and more generally human-computer interaction. One potentially beneficial application of this system, in particular the vibrotactile feedback device, would be the use of vibrotactile feedback for motor rehabilitation for patients recovering from stroke or other neurological injury. Such a vibrotactile feedback device could provide a low-cost, portable solution to help retrain motor skills. In fact, such a device could be tailored for at-home or even chronic use in disabled patients, as opposed to the conventional physical therapy sessions at a rehabilitation clinical.

In summary, the goal of this work was to develop the technological framework that would not only further improve the technical capabilities of the research community, but would also provide insight into how state-of-the-art technology could be implemented for scientific investigation in the future.

APPENDIX A

CORRELATION BETWEEN REACH DURATION AND PERFORMANCE

As discussed in section 4.2.2, subjects tended to make slower reaches when using vibrotactile feedback. This raised the question as to whether the improvement in performance was simply due to the fact that the tactile feedback forced subjects to slow down and pay more attention, or if the vibrotactile feedback actually served as an informative sensory signal. To address this issue, a correlation analysis was performed between reach duration and movement error. To characterize this relationship, correlation coefficients were calculated for each task condition for each individual subject (as shown in Figure 4.6). Figure A1 below shows scatter plots of reach duration against movement error grouped by task condition. Linear fits to the raw data are shown for each of the 9 subjects (red lines) on each of the 8 plots representing the 8 task conditions.
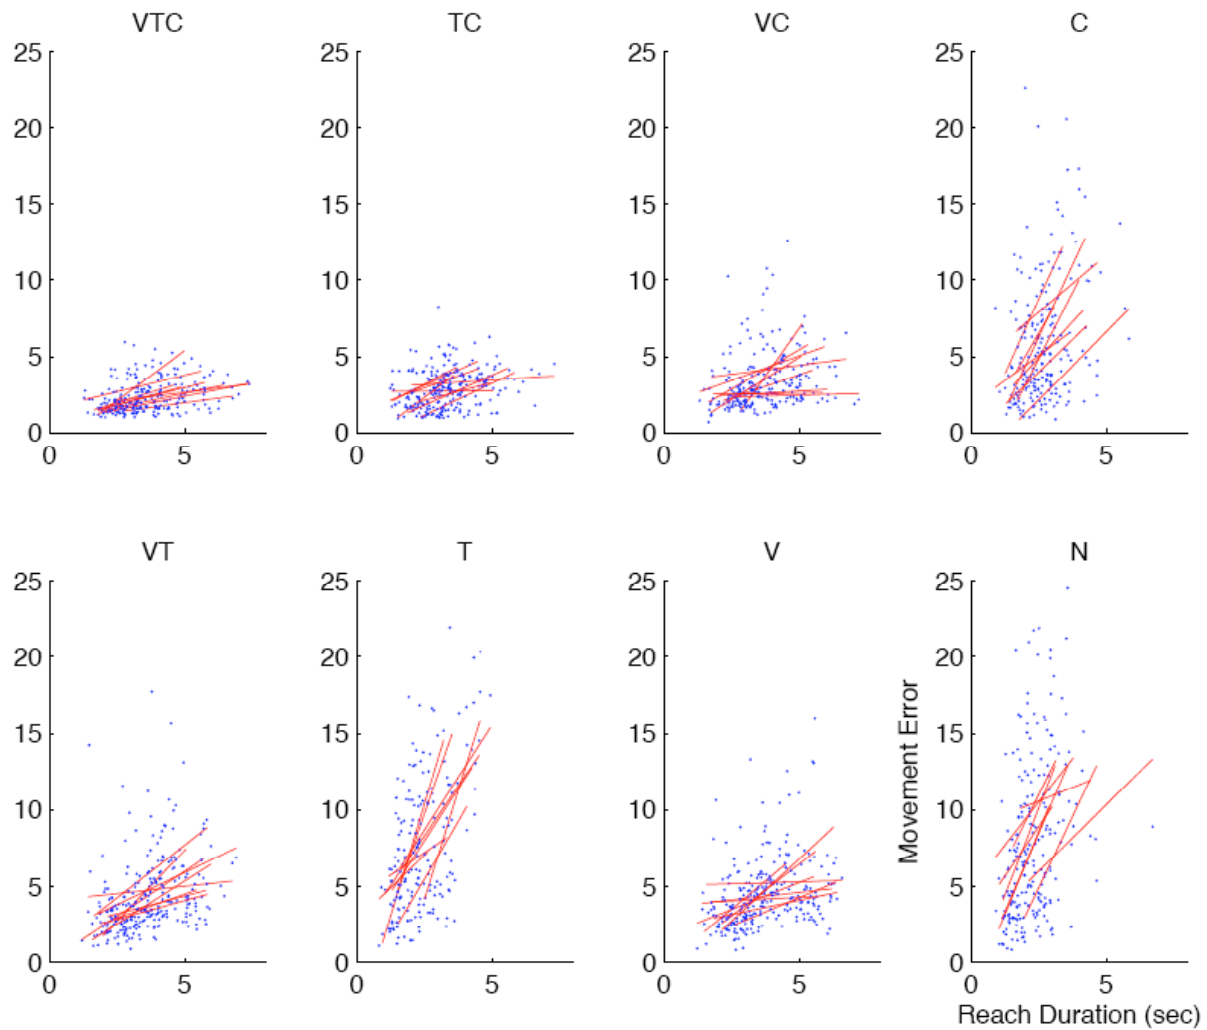
Figure A 1: Relationship between reach duration and movement error. Scatter plots show data points (blue dots) and linear fits (red lines) for all subjects (n = 9) within each task condition. V = vibrotactile on, T = trajectory visible, C = cursor visible, N = none.

As is shown in Figure A1 above, subjects tended to show a positive correlation (i.e. positive slope to the linear fits) between reach duration and movement error, demonstrating that subjects actually tend to do worse at the task when moving slower for a given task condition.

BIBLIOGRAPHY

ACA. Action Plan for People with Limb Loss. In, vol 2005. ACA (2005).

ALSA. Facts you should know about ALS. ALS Association. <http://www.alsa.org>. (2007).

Bach-y-Rita P, Collins CC, Saunders FA, White B, and Scadden L. Vision substitution by tactile image projection. Nature 221: 963-964 (1969).

Bild DE, Selby JV, Sinnock P, Browner WS, Braveman P, Showstack JA. Lower extremity amputation in people with diabetes. Epidemiology and prevention. Diabetes Care. 12: 24-31 (1989).

Bryant CL, Gandhi NJ, Real-time data acquisition and control system for the measurement of motor and neural data. J Neurosci Methods 142 (2), 193-200 (2005).

Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, and Nicolelis MA. Learning to control a brain-machine interface for reaching and grasping by primates. PLoS Biol 1: E42 (2003).

Center NSCIS. Spinal Cord Injury: Facts and figures at a glance. In. National Spinal Cord Injury Statistical Center (2005).

Chatterjee A, Aggarwal V, Ramos A, Acharya S, and Thakor NV. A brain-computer interface with vibrotactile biofeedback for haptic information. J Neuroeng Rehabil 4: 40 (2007).

Cincotti F, Kauhanen L, Aloise F, Palomaki T, Caporusso N, Jylanki P, Mattia D, Babiloni F, Vanacker G, Nuttin M, Marciani MG, and Del RMJ. Vibrotactile feedback for brain-computer interface operation. Comput Intell Neurosci 48937 (2007).

Debus T, Jang TJ, Dupont P, Howe R. Multi-channel vibrotactile display for teleoperated assembly. In: Proceedings IEEE International Conference on Robotics and Automation. (2002).

Dhillon GS, Horch KW. Direct neural sensory feedback and control of a prosthetic arm. IEEE Proc IEEE Eng Med Biol Soc. **7**: 5329-32 (2005).

Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, and Donoghue JP. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442(7099), 164-171 (2006).

Jones LA, Lockyer B, Piateski E. Tactile Display and Vibrotactile Pattern Recognition on the Torso. Advanced Robotics, Vol. 20, No. 12, pp. 1359 – 1374 (2006).

Kaczmarek KA, Webster JG, Bach-y-Rita P, and Tompkins WJ. Electrotactile and vibrotactile displays for sensory substitution systems. IEEE Trans Biomed Eng 38: 1-16 (1991).

Kim YK, Yang X. Hand-writing Rehabilitation in the Haptic Virtual Environment. IEEE Workshops on Haptic Audio Visual Environments and their Applications. Ottawa, Canada 4-5 November (2006).

Kuiken TA, Li G, Lock BA, Lipschutz RD, Miller LA, Stubblefield KA, Englehart KB. Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms. JAMA. 301(6):670-1. (2009).

Lieberman J, Breazeal C. TIKL: Development of a wearable vibrotactile feedback suit for improved human motor learning. IEEE Transactions on Robotics 23 (5), 919–926. (2007).

London BM, Jordan LR, Jackson CR, Miller LE. Electrical stimulation of the proprioceptive cortex (area 3a) used to instruct a behaving monkey. IEEE Trans Neural Syst Rehabil Eng 16 (1), 32-36 (2008).

Mackenzie IS, Kauppinen T, Silfverberg M. Accuracy measures for evaluating computer pointing devices. Proc. the SIGCHI Conf. on Human Factors in Computing Systems (CHI'01, Seattle, WA, United States) pp 9–16 (2001).

Marasco PD, Schultz AE, Kuiken TA. Sensory capacity of reinnervated skin after redirection of amputated upper limb nerves to the chest. Brain. 132(Pt 6):1441-8. (2009).

Marston JR, Loomis JM, Klatzky RL, Golledge RG. Nonvisual route following with guidance from a simple haptic or auditory display. Journal of Visual Impairment & Blindness, 203-211. (2007).

Nagel SK, Carl C, Kringe T, Martin R, and Konig P. Beyond sensory substitution--learning the sixth sense. J Neural Eng 2: R13-26, (2005).

Ojakangas CL, Shaikhouni A, Friehs GM, Caplan AH, Serruya MD, Saleh M, Morris DS, Donoghue JP. Decoding movement intent from human premotor cortex neurons for neural prosthetic applications. J Clin Neurophysiol 23: 577-584, (2006).

Reed CM, Delhorne LA. Current results of a field study of adult users of tactile aids. Seminars in Hearing 16, 305–315. (1995).

Robinson DA. A Method of Measuring Eye Movement Using a Scleral Search Coil in a Magnetic Field. IEEE Trans Biomed Eng 10, 137-145 (1963).

Romo R, Hernandez A, Zainos A, Brody CD, and Lemus L. Sensing without touching: psychophysical performance based on cortical microstimulation. Neuron 26: 273-278, 2000.

Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. A high-performance brain-computer interface. Nature 442 (7099), 195-198 (2006).

Schwartz AB, Cui XT, Weber DJ, Moran DW. Brain-controlled interfaces: movement restoration with neural prosthetics. Neuron 52:205-20 (2006).

Scott SH, Loeb GE. The computation of position sense from spindles in mono- and multiarticular muscles. Journal of Neuroscience, 14, 7529–7540 (1994).

Taylor DM, Tillery SI, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. Science 296 (5574), 1829-1832 (2002).

van Beers RJ, Wolpert DM & Haggard P. When Feeling Is More Important Than Seeing in Sensorimotor Adaptation. Current Biology, Vol. 12, 834–837, (2002).

van Erp JBF, van Veen HAHC. A multi-purpose tactile vest for astronauts in the international space station. In: Proceedings of Eurohaptics. (2003).

Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. Nature 453: 1098-1101 (2008).

Vuillerme N, Chenu O, Pinsault N, Moreau-Gaudry A, Fleury A, Demongeot J, and Payan Y. Pressure sensor-based tongue-placed electrotactile biofeedback for balance improvement-- biomedical application to prevent pressure sores formation and falls. Conf Proc IEEE Eng Med Biol Soc 2007: 6114-6117, (2007).

Wheatstone C. Contributions to the physiology of vision - Part the first; On some remarkable and hitherto unobserved phenomena of binocular vision. Philosophical Transactions of the Royal Society of London, 128, 371-394 (1838).

Wheatstone C. The Bakerian Lecture: Contributions to the physiology of vision - Part the second; On some remarkable and hitherto unobserved phenomena of binocular vision. Philosophical Transactions of the Royal Society of London, 142, 1-17 (1852).