TWuiST: A Discrete Tactile-Proprioceptive Display for Eye and Ear Free Output on Mobile Devices

Tony Morelli* and Eelke Folmer†

Department of Computer Science and Engineering
University of Nevada, Reno

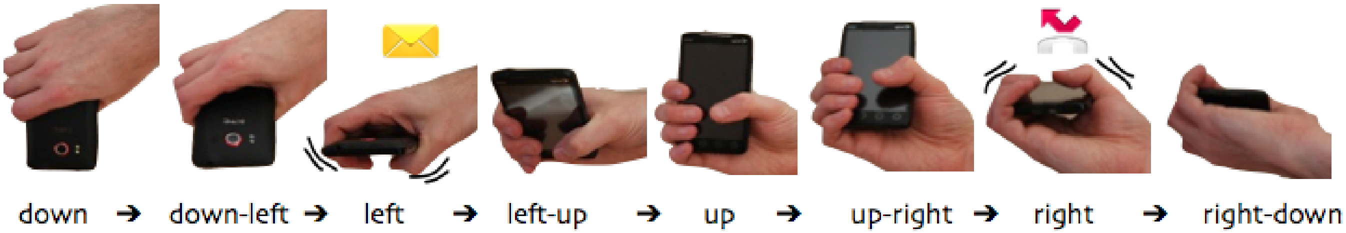

Figure 1: The user scans through eight different orientations of their mobile device where a specific target orientation is indicated using a vibrotactile cue. Different target orientations correspond to different types of information that the user can access in an ear and eye free manner. An example application of this technique indicates the status of a mobile device where LEFT indicates a new email and RIGHT a missed call.

ABSTRACT

Proprioception –the human ability to sense the orientation of limbs without vision or hearing– is one of the main drivers of complex motor operation, which is something mobile interfaces may be able to exploit to achieve robust eye and ear free forms of interaction. This paper explores the use of proprioception as an output modality by combining kinesthetic information of a mobile device with vibrotactile feedback. A user study with 16 users explored the temporal resolution of proprioceptive displays for two different spatial resolutions with orientations (o) either defined in a space ($o = 6$) or in a plane ($o = 8$). Users were able to find target orientations in $2,466$ (space) and $2,612$ (plane) milliseconds. The performance of discrete proprioceptive displays is comparable with more advanced forms of tactile feedback provision -but unlike these- proprioceptive displays can be facilitated using features already present in current mobile devices. Our experiences elicited a number of guidelines and tradeoffs for the design of discrete proprioceptive displays.

Index Terms: H.5.2 [HCI]: User Interfaces—Haptic I/O

1 INTRODUCTION

Interaction capabilities of mobile devices are typically restricted by their weight and size, which curbs their functionality, usability, and accessibility. For example, mobile devices increasingly feature touch screens to optimize available input and output space, but interacting with onscreen keyboards often proves to be error prone due to their small buttons. Their lack of tactile feedback is also an important barrier towards their use by users with visual impairments.

Alternative interaction techniques have been developed that seek to increase input space of mobile devices without compromising their portability. Many of these techniques seek to appropriate the device itself into an input device by: (a) sensing its orientation [1, 2]; (b) sensing its position on a flat surface [3]; and (c) sensing of gestures made with the device [4, 5] or against the device [6].

Whereas most research has focused on increasing available *input* space, but we argue that *output* capabilities of mobile devices are also constrained, not only because of their small screens and limited audio capabilities, but because: (1) users often interact with their mobile devices when they are active and the use of a display or audio may impede the users' safety, for example, when they are walking or driving; (2) the use of audio feedback may be limited due to noisy environments and safety and privacy issues; and (3) users may be unable to use a screen because of a visual impairment.

Tactile feedback lends itself well to the design of mobile interfaces to achieve eye and ear free interaction [7] as mobile devices are often held closely to our skin. However, tactile feedback provision on mobile devices remains relatively underutilized [8, 9]. Mobile devices typically feature a single rotary mass motor that provides on/off vibrotactile feedback with a fixed frequency and their latency limits the use of sophisticated drive signals. Consequently haptic feedback on mobile devices is restricted to tactons that vary only in rhythm or pulse duration [9].

Proprioception is an interoceptive (internal stimuli) sensory modality that is distinct from exteroceptive (external stimuli) sensory modalities such as sight, touch, and hearing. Proprioception allows humans to sense the position and orientation of their limbs without exteroception [10]. Being one of the main drivers for complex motor operations [11], the role that proprioception plays in input modalities such as gestures and touch, is self-evident and significant. An exteroceptive modality, such as haptic feedback, can be augmented with proprioceptive information to facilitate a significantly larger information space that can be accessed in an ear and eye free manner. For example, *tactile-proprioceptive displays* [12, 13] have been explored in mobile navigation systems to point out an object of interest in the user's vicinity using the user's own arm (see related work following for a comprehensive overview).

Rather than using different directions an arm can point at as indicated using vibrotactile feedback, we explore the use of different orientations of a mobile device itself to convey information using semaphores that can be retrieved in an ear and eye free manner. Specifically we seek to convey information to the user that is not constrained to relate to the user's immediate physical environment. Our technique could be used to retrieve a mobile device's status while driving or to increase tactile output capabilities of mobile devices as to make them more accessible to users who are visually impaired.

---

*e-mail: morelli@cse.unr.edu
†e-mail: efolmer@cse.unr.edu

2 BACKGROUND AND RELATED WORK

In performing a motor operation towards a target, humans generally combine visual and proprioceptive information; but if the location of a target has been memorized, proprioceptive feedback alone is sufficient for relating the position of the users' limb to the target. Consequently, proprioception has been explored for developing robust ear and eye free forms of *input* for mobile devices [2, 14] and virtual environments [15]. Vibrotactile feedback has been explored in mobile tilting interfaces [7] where different tactile patterns are used to convey the speed with which data scrolls. Wrist tilt has been evaluated for input provision on mobile devices [16]. Leveraging proprioception to indicate targets on a single axis as an *output/display* technique was recently explored in the following approaches.

Sweep-Shake [13] is a mobile phone application that can point out geo-located information by sweeping the phone in the user's environment. The phone's compass and GPS are used to determine the user's location and the direction in which the phone is pointing. Directional vibrotactile feedback (increasing its magnitude) conveys the location of an object of interest. Authors do not provide information about the size of the targets or the window around the target upon which vibrotactile feedback is provided. Users can perform gestures to interact with the object of interest but these are not directed at the acquired target. A study with four sighted users found they were able to find targets in a 360° circle around them on average in 16.5 seconds.

Ahmaniemi [12] explored finding targets using a mobile device that consists of a high precision inertial tracker (gyroscope, compass and accelerometer) and a C2 vibrotactor. Two types of vibrotactile cues were explored for rendering targets: (1) an on target cue (260Hz sine wave mixed with a 30Hz envelope signal); and (2) a directional cue using a tactile window of 10° around the target (using the same on target cue where the frequency and amplitude of the envelope shape were increased linearly). Targets were rendered randomly at eight different locations on a 90° horizontal line with varying widths. A user study with eight sighted users found they were able to find targets on average in 1.8 seconds. No significant difference was found between vibrotactile feedback provision for efficiency and target size, though smaller targets took longer to find than larger targets. Target sizes larger than 15° were most effective. Directional vibrotactile cues are more efficient than non-directional cues when target distance is furthest but it negatively affects finding targets that are close. It makes it also harder to distinguish targets that are close to each other as distinguishing the edges of a target becomes harder.

VI Bowling [17] is a tactile/audio exergame for users who are blind that is played using a commercial-of-the-shelf motion sensing controller with an integrated vibrotactor (Wii remote). A technique called "tactile dowsing" guides the player to point their controller at the pins using directional vibrotactile feedback, where the position of the controller is tracked using a wireless IR emitter peripheral. Proprioception is used to facilitate a basic form of motor learning as users throw a virtual bowling ball at the sensed location of the pins. With a close-to-target window of 38.6° and a target size of 7.2° a user study with six legally blind adults found that users were able to find the target on average in 8.78 seconds and users were able to perform a directed gesture (throw) with an average error of 9.76° towards the target.

Virtual shelves [2] is an input technique where users triggers shortcuts by orienting a Wii remote within the circular hemisphere in front of them. This technique includes a training phase where vibrotactile feedback is provided to identify targets and to facilitate the development of a spatial memory. Authors refrain from specifying whether directional or non-directional feedback is provided to scan to a target in 2D but given the technical limitations of a Wii remote [17] to only provide directional vibrotactile feedback using the modulation of pulse length, directional feedback using a wii remote is restricted to indicate a target on a single axis.

Magnusson [18] evaluates a system similar to Sweep-Shake [13], where a non-directional audio cue indicates whether the user is pointing the phone within a window that contains a beacon that users must physically approach. The phone's location and the direction it is pointing are acquired using GPS and a compass. Vibrotactile feedback is increased when the user gets closer to the beacon. Different sized target windows are evaluated with 15 sighted users, where a window of size 30° to 60° was found to be most efficient. PointNav [19] is an extension of the Magnusson system but modified as to provide a 50ms non-directional vibrotactile cue when the phone points within a 30° window of the object of interest. PointNav was evaluated qualitatively with five visually impaired users where all users were able to find all beacons, but no performance data is provided.

All previous approaches have predominantly explored proprioceptive displays for 1D target acquisition where a target typically refers to a real world object. As such these displays are "analog" in nature as they represent a continuous value (such as the direction to a point of interest). In this work we take a different approach where we explore tactile-proprioceptive displays for conveying "discrete" information, e.g., bits of information that do not relate to the user's immediate context, which allows for proprioceptive displays to be more flexible.

3 PROPRIOCEPTIVE DISPLAYS

Proprioception is the human ability to sense the position and orientation of the limbs and its extremities without using vision [10]. We define a *proprioceptive display* to be an output technique that turns the space of allowable and distinguishable positions and orientations of a limb into an information space that the user has access to in an ear and eye free manner using proprioception. Prior to having access to this information space, the user orchestrates their limb into a specific orientation or position as indicated using an exteroceptive stimulus. Proprioceptive displays should be considered "hybrid" displays that rely upon a small amount of feedback in another modality for the arrangement of the limb as to facilitate a significantly larger information space.

All techniques discussed in the related work section implement proprioceptive displays using a handheld orientation aware mobile device, for example, a mobile phone with integrated compass. Initially when the user holds the mobile device in their hand, the direction it is pointing is sensed using *stereognosis*, e.g., the perception of the shape of a 3D object using touch. Mobile devices are typically elongated and have identifiable tactile features, such as buttons, which further facilitates this process. Users scan their environment either by adjusting the direction their forearm [12, 13] or stretched arm [17] is pointing. A vibrotactile cue indicates when the mobile device is pointed at the target, upon which the direction of the target is conveyed to the user using proprioception using the current position of their arm or forearm. Though a limb could be orchestrated using sonification, haptic feedback facilitates ear and eye free interaction and is readily available on most mobile devices.

3.1 Information Space

The information space of a proprioceptive display is constrained by the amount of available rotation of a limb. Adjusting the direction in which the arm or forearm is pointing involve rotations of the shoulder (Glenohumeral) joint. With approximately 135° of flexion, 40° of extension, 170° of abduction and 35° of adduction, the shoulder joint is the most mobile joint in the body [20]. Flexion, extension, abduction and adduction are involved in adjusting the direction the arm is pointing and abduction and adduction determine the direction the forearm is pointing. The arm or forearm can freely rotate within in a hemisphere defined in front or to the side of the user, which creates an information space that is able to point out a 3D vector relative to the user. Techniques discussed in the related work section have only explored a subset of this space as the user scans a horizontal line to find a target, only conveying a 2D vector.

3.2 Spatial Resolution

Proprioceptive displays have a relative large spatial resolution, e.g., a previous study found users were able to detect target sizes as small as 5° for a 90° window [12]. Given the small amount of haptic feedback, i.e, a single vibrotactor that is used for orchestrating the limb this is a significant increase in the *type* of information that can be communicated to a user. A proprioceptive display can convey three "analog" values (a 3D vector) simultaneously, something that is difficult to achieve with available haptic capabilities of current mobile devices. Concomitantly, the temporal resolution of proprioceptive displays is small as users can only acquire the target by searching the information space. The generated information space has been explored for conveying information such as: (1) a point of interest [13]; (2) a target to hit [17]; or (3) a direction to follow [18]. To determine the orientation these displays require an external point of reference. Using a compass may yield conflicting data in indoor environments and tracking the position of a mobile devices using a camera [17] is not feasible in mobile contexts.

3.3 A Discrete Proprioceptive Display

We explore the use of a proprioceptive display for conveying information that is not constrained to relate to the user's immediate physical or virtual environment. Specifically, we explore a *discrete* proprioceptive display to communicate bits of information using semaphores. Rather than using different directions an arm can point in, we explore the use of different orientations of the mobile device itself. The shoulder joint is constrained to rotate within a hemisphere –whereas the hand can point to almost any direction in a sphere through a combination of: (1) rotations of the wrist (60° of flexion, 45° of extension, 25° of abduction, 65° adduction); (2) rotations (80° of pronation, 88° of supination) of the forearm at one of the elbow (proximal radioulnar) joints; and (3) rotations of the shoulder joint, 170° of abduction and 35° of adduction [20]. Figure 2 shows types of rotations of the wrist and the forearm. Implementing a discrete proprioceptive display using hand orientations allows for creating an information space that is twice the size of that of the arm and as the hand can be adjusted using smaller motions than the arm, users may be able to find targets faster.

3.4 Auto Semaphoring

Unlike analog proprioceptive displays, discrete proprioceptive displays use distinct predefined orientations of the mobile device, which –when the user has successfully memorized these– can be scanned through more efficiently as users may be recall the orientations from kinesthetic memory. A vibrotactile cue indicates the target orientation, which corresponds to a unique hand orientation, which the user has access to in an ear an eye free manner. The found hand orientation can then be used for communicating "discrete" information using semaphores. This concept –and a possible application for discrete proprioceptive displays– is best illustrated using the following scenario. Mobile devices typically alert users to an internal state change (new text message, missed call, new voicemail, new email, low battery) through the same vibrotactile cue. The user has to look at their screen to retrieve the state change. Using a discrete proprioceptive display a state change can be retrieved while leaving ears and eyes free for sensing the immediate environment, for example when driving a car. Upon receiving the alert, the user takes their phone and scans through the predefined orientations to find the target orientation that is indicated using a vibrotactile cue. Distinct orientations of the hand (i.e. semaphores) correspond to different states, for example, the phone pointing left indicates a new voicemail and pointing right indicates a missed call (see Figure 1). Looking at a mobile screen while driving impedes safety, however, proficient drivers can shift gears manually without taking their eyes off the road [2]. The motor operations required for scanning through the different orientations require little cognitive load and may not impede the ability to drive.

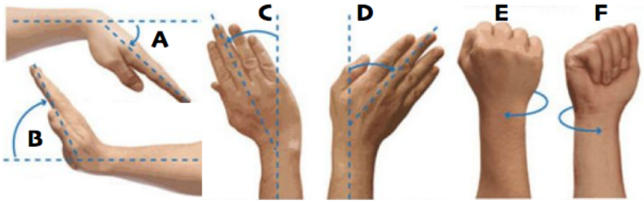

Figure 2: Available wrist rotations: 60° of flexion (A) 45° of extension (B) 25° of abduction (C) 65° of adduction (D) 80° of pronation (E) 88° of supination (F).

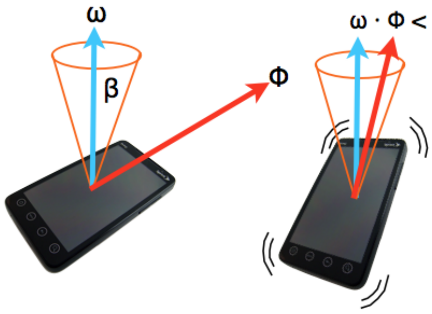

Figure 3: The orientation $\boldsymbol{\Phi}$ of the mobile device is outside the tactile window β defined around ω (left) and inside the tactile window β upon which vibrotactile feedback is provided (right).

3.5 Formal Description

For proprioceptive displays in general, an initial orientation of a mobile device can be formally specified as $\boldsymbol{\Phi} = (\phi, \theta, \omega)^T$ where ϕ is roll, θ is pitch, and ω is yaw. Because mobile devices are elongated their distinct shape aligns with $\boldsymbol{\Phi}$, whose direction the user perceives through stereognosis. A target orientation can be specified as: $\boldsymbol{\omega} = (u, v, w)^T$ and a window of β degrees is defined around the target orientation. Users change the orientation of their mobile device by adjusting the orientations of their hand until the angle between the current and target orientation is smaller than β. When $\boldsymbol{\omega} \cdot \boldsymbol{\Phi} < \beta$ a vibrotactile cue is provided (see Figure 3) to indicate to the user that the current orientation matches the target orientation. Haptic feedback is provided to facilitate ear and eye free interaction, which is useful in mobile contexts. Because the user scans through predefined orientations, no directional feedback is required to indicate a target orientation.

3.6 Implementation

We use a commercially available smart phone (HTC Evo) to demonstrate that a discrete proprioceptive display can be implemented using its existing features. Device orientation is determined as follows. Mobile devices increasingly feature accelerometers [6], which measure linear acceleration of the device along three axes. The accelerometer's axes are aligned with the mobile device's chassis and because the earth's gravity is a force that causes an acceleration, this can be used for determining the device's orientation, for example, to adjust the screen orientation. Accelerometers are limited in their ability to report the exact orientation of the mobile device. When the mobile device is positioned flat on a table, rotations about its length (roll) and width (pitch) can be retrieved using

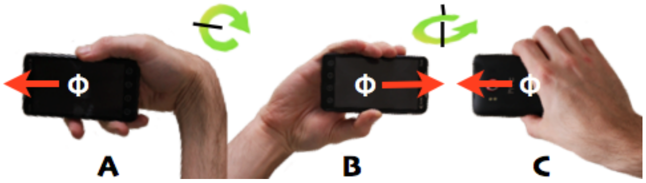

Figure 4: Rotations around the center of the phone (A→B) do not allow for discriminating different values of Φ but including a rotation around its width (C→B) generates two different values for pitch. Photos taken from a top-down view with the user facing forward.

its accelerometer, but not any rotations around its center (yaw). Yaw can be determined by combining accelerometer data with either a magnetometer or a gyroscope. Magnetometers may be subject to interference in the presence of large metal objects whereas gyroscopes have only recently become available in smart phones.

3.7 Designing an Information Space

We design our information space based on what orientations can be feasibly measured using a single accelerometer. Whether the mobile device is pointing up or down can be determined using pitch, but determining unique values for Φ without yaw is impossible when Φ lies in the horizontal plane, e.g., it is not possible to differentiate between the mobile device pointing to the left, forward, to the right or backwards when the phone is rotated around its center. However, two unique orientations of the phone in the horizontal plane can be determined by including a rotation about the width of the phone for a single pair of opposing directions. For example, rather than pointing the mobile device towards the left or the right using abduction and adduction of the wrist joint, we can "twist" the mobile device over its width using pronation and supination of the forearm, where different values for pitch allow for discriminating two orientations in the horizontal plane (See Figure 4). Using a twist over either the phone's width or length, changes in roll and pitch allow for identifying six unique orientations that all involve supination and pronation of the forearm. Specifically, a vector indicating earth's gravity can be retrieved from the phone's 3-axis accelerometer as it is decomposed along each axis. When an axis is horizontal this value is 0 and when vertical -10.0 or 10.0. When the value on any axis is maximal or minimal, six distinct orientations ($o = 6$) can be defined in a space ($n = 3$). Four orientations involve supination/pronation, e.g., [RIGHT, DOWN, LEFT, UP] and two involve flexion/extension, e.g., [BACK, FORWARD]. Orientations are shown in Figure 5.

A problem with this set is that we can't exploit the elongated features of the phone to point it in a particular direction as the orientations [BACK, FORWARD] have the phone pointing in the same direction (see Figure 5). To circumvent this limitation we defined a mnemonic that helps the user recall each distinct orientation by having them visualize the phone to be held inside a cube where the screen of phone must face each of the inside faces of the cube to cycle through the different orientations as shown in Figure 6:left. Though in-between orientations [BACK-UP] could be measured, we restrict ourselves to six orientations for the cube mnemonic to work.

3.8 Optimizations for Scanning

Six unique orientations defined in a space ($o = 6$) creates an information space with a size of $\log_2(6) = 2.58$ bits. These orientations fully utilize all possible rotations of the hand and are most distinguishable from each other, as any pair of orientations can be expressed as vectors (u, v) that are either orthogonal ($u \cdot v = 0$) or each others inverse ($u = -v$). Preliminary experiences with this information space revealed that the time it takes to find a random orientation is significantly affected by how a user scans through these orientations. Ahmaniemi [12] found in their study that users could find a random target on average in 1.8 seconds but their targets were all located on a horizontal 90° line and the user could scan to each target using a constant velocity. The shortest path between orientations defined in a space (that do not lie in a plane) cannot be scanned through with a constant velocity.

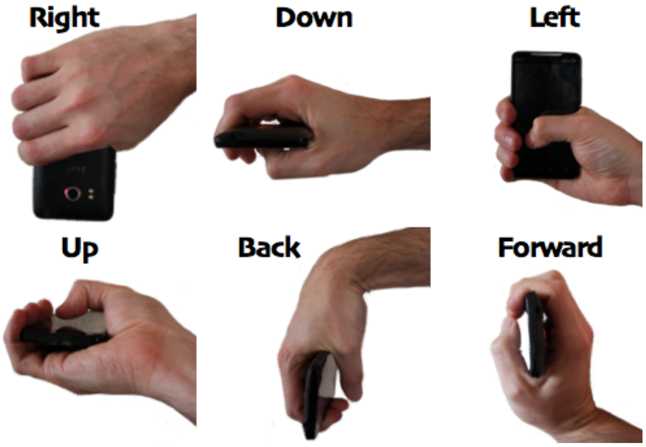

Figure 5: Six disjunct orientations defined in a space. User is facing to the left, photos taken from the side.

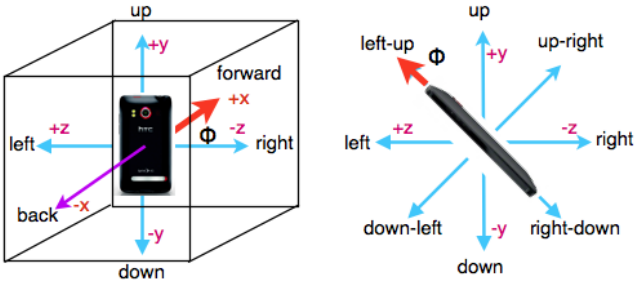

Figure 6: Information space for six orientations in a space using the cube mnemoic (left) and for eight orientations in a plane (right).

The largest range of motion (ROM) of the hand in a plane (∼373°) can be achieved through supination (88°) and pronation (80°) of the forearm supplemented with adduction (35°) and abduction (170°) of the shoulder [20]. This rotation starts in the pronated position of the forearm with the shoulder adducted to about 90°, followed by a combined supine rotation of the forearm and abduction of the shoulder joint. A subset of orientations can be identified [RIGHT→ DOWN→ LEFT→ UP] from our set that can be scanned through with constant velocity using a 270° rotation around the pronation/supination axis (phone's width) where the user starts scanning in either RIGHT (clockwise) or LEFT (counterclockwise) orientation. Scanning the last orientations [BACK,FORWARD] takes more time as they involve larger or combined rotations. From UP a combined 90° of pronation (forearm) with extension (wrist) is required to get to BACK and 180° of flexion (wrist) to get to FORWARD.

To optimize for scanning and to optimally explore the spatial resolution of discrete proprioceptive displays an alternative information space was defined that consists of eight equidistant orientations in a plane ($o = 8$), which can be scanned through in a clockwise 315° rotation using a constant velocity. This plane is orthogonal to the pronation/supination axis. The spatial resolution of wrist-tilt interaction for mobile input has been previously investigated [16, 21]. Rahman [16] found that users can distinguish up to 16 levels on the pronation/supination axis but this type of input is coupled to visual feedback. Because we rely upon proprioceptive feedback and not on visual feedback only 8 orientations are explored which are defined 45° degrees apart from each other, which may be easy to recall for users. This space has a size of $\log_2(8) = 3$ bits. A benefit of this space is however is that the elongated features of the phone can be used to point the phone in eight different directions, e.g., [DOWN, DOWN-LEFT, LEFT, LEFT-UP, UP, UP-RIGHT, RIGHT, RIGHT-DOWN] (see Figure 1) which may be easier for the user to use than recalling orientations using the cube mnemonic. Figure 6 shows both information spaces. Orientations with identical names in both sets are transposed 90° from each other.

4 USER STUDIES

We evaluated the temporal resolution of both information spaces using a game (TWUIST-N-LOCK) that was developed in Java using the Android SDK. Because games are considered powerful motivators, this may allow for measuring optimal performance. The goal of the game is to find as many target orientations within a time limit.

4.1 Instrumentation

A target orientation is indicated using vibrotactile feedback. The HTC Evo features a 4mm single rotary mass motor that provides on/off vibrotactile feedback with a frequency of 240Hz and an amplitude of 0.5G [1]. A tactile window in which vibrotactile feedback is provided was defined as follows. Ahmaniemi [12] found target sizes larger than 15° to be most effective for an analog proprioceptive display. Based on early trials, a target size of 18° was found feasible, when the value on the target axis is within 2.0 of either the minimum (-10) or maximum (10) the phone's vibrotactor is activated. For $o = 8$, orientations such as DOWN-LEFT that are defined between the axes, values for both Y and Z axes must be within 1.0 of half of the minimum or maximum value resulting in the same 18° target size. Trials and previous studies [12, 17] found no significant interference of vibrotactile feedback with accelerometer readings when the phone is handheld. A potential risk with using accelerometers is fuzzy input, though Kalman filters can be used to increase their accuracy. Because we use fixed orientations and not random orientations, trials also revealed that neither a Kalman filter nor a more precise accelerometer was required to be able to measure these orientations. A target orientation is determined at random, as it can be assumed that the type of message to be communicated using auto-semaphoring is not known a priori.

To optimize for scanning the game requires the user to start scanning in the start position, for $o = 6$ this is RIGHT and for $o = 8$ this is DOWN. These are essentially the same orientations (see Figure:1 and 5) but named differently for each information space to facilitate the user recalling these orientations. Because each search starts from the initial position, this position is excluded from the set of target orientations and users are effectively scanning for only five ($o = 6$) or seven ($o = 8$) random target orientations. The start orientation becomes the target orientation after the user finds a random orientation. This is useful as it allows for measuring how fast the start orientation can be assumed as the phone's initial position is not a priori known before the user uses the discrete proprioceptive display. Users are required to hold the phone in the tactile window for at least 1 second until the sound of a bell is heard. This was implemented to avoid generating false positives, which happen when the phone is moved too vigorously. After finding the target orientation, the user then goes back to the start position, which is indicated the same as a target orientation but instead of a bell the sound of a drum is played, to prompt the player to start their search. A score computed after finding each random target, which is $score = 20 - t$ with t the number of seconds it took to find the target orientation. This score is announced at the end of the game.

4.2 Participants

We recruited sixteen users (3 female, 13 male, average age 28.4, SD=4.67) to participate in a user study. All were right handed and none had any impairments in tactile perception or motor control. Fourteen users owned a smartphone and two owned a regular phone. Participants were randomly divided into two groups (A,B) with 8 participants in each group. Group A played the game using ($o = 6$) and group B using ($o = 8$). Participants in group A were taught all six orientations using the "cube" mnemonic. Participants were instructed to start each search in the RIGHT orientation and then scan through the orientations in the following order: DOWN→ LEFT→ UP→ BACK→ FORWARD (See Figure 5). Participants in group B were taught all eight orientations using the elongated features of the phone to point it in a direction. They were instructed to start in DOWN and then scan the sequence: DOWN-LEFT→ LEFT→ LEFT-UP→ UP→ UP-RIGHT→ RIGHT→ RIGHT-DOWN (see Figure 1). Prior to playing the game, participants were asked to scan through the sequence of orientations. None of the participants expressed experiencing any discomfort in performing the rotations.

4.3 Procedure

All participants were informed they would have to find as many orientations in four minutes as possible and that the faster the correct orientation was found, the higher their score would be. Participants were instructed to scan for orientations with the back of the phone pressed against the palm of the hand (see Figure 1 & 5) and not rotate the phone with their fingers. Participants were instructed upon finding the target orientation to return to the down position and only to start scanning after hearing the drum audio cue. User studies took place in a small office room. Participants played the game using their dominant hand while seated on a chair. An observer was present for all user studies. To avoid interference, the screensaver and the auto rotate features on the phone were disabled and the buttons at the bottom of the screen were covered with tape. Nothing was displayed on the screen. Participants received a $5 gift certificate for their participation. A timer was displayed on another phone in front of the user on a desk indicating the amount of time left. The game would automatically stop after 240 seconds and record the following data to a log file: (1) each target orientation; (2) orientations scanned through prior to finding the target orientations where we sampled and averaged orientations in 100 millisecond intervals; and (3) the time it took to find the target orientation (search time) from which the 1-second holding time is subtracted.

4.4 RESULTS

Table 1 and 2 list the results of the user study for both information spaces for the four minutes of playing time. Column 1 lists the target orientation and for comparison, identical orientations in both spaces, e.g., RIGHT and DOWN have the same number behind their name. For ($o = 6$) users found a total of 655 orientations (M=81.9, σ=18.3) and for ($o = 8$) users found 634 orientations (M=79.3, σ=13.0). The variance of search times for all orientations is large, i.e., σ^2= 330 seconds for $o = 6$ and σ^2=248 seconds for $o = 8$, which is attributed to probability as target orientations were picked at random. Our analysis henceforth focuses on performance data of each orientation. Table 1 and 2 column three lists the average search time for each orientation and column four lists the standard deviation. Figure 7 shows a combined graph of the average search times for the sequence of orientations that the user scans through for both information spaces. As anticipated, for ($o = 6$) the average search times for BACK and FORWARD are higher than their expected values. A linear regression line has been plotted based on the first four orientations and which is extrapolated to show their predicted values for the last pair if these orientations could be scanned to with a constant velocity. A visual analysis and a Grubbs' test using consecutive differences shows that the average search time for BACK is

---

[1] The exact specifications of the HTC Evo's pager engine is considered proprietary information, therefore the specifications of a third party replacement pager engine for this phone are used, which can be assumed is similar.

Table 1: Results for six orientations.

Orientation	**Frequency**	**Time** (ms)	**Stdev** (ms)
RIGHT (1)	327	1,100	875
DOWN (2)	70	1,576	2,559
LEFT (3)	69	2,094	1,137
UP (4)	65	2,363	1,326
BACK	62	3,689	2,528
FORWARD	62	3,975	1,683
Total	**655**	**2,466**	**1,147 (avg)**

Table 2: Results for eight orientations.

Orientation	**Frequency**	**Time** (ms)	**Stdev** (ms)
DOWN (1)	317	1,058	740
DOWN-LEFT	43	3,616	2,713
LEFT (2)	52	1,707	921
LEFT-UP	46	2,334	879
UP (3)	47	2,375	1,267
UP-RIGHT	40	2,809	1,238
RIGHT (4)	41	3,252	1,890
RIGHT-DOWN	48	3,744	2,618
Total	**634**	**2,612**	**934 (avg)**

a statistically significant outlier ($\boldsymbol{Z} = 1.73, N = 6, p < .05$).

Because a tactile window of 18° was implemented, average scanning velocities are approximated based on their average search times. The first four orientations are scanned through with 114°/second and the last two with 68°/second and 633°/second. The scanning velocity to BACK is lower than for the first four orientations, but to FORWARD it is higher. This last result is somewhat counter intuitive as it involves a 180° rotation. However, this rotation involves a smaller motion, e.g., extension of the wrist as opposed to a rotation of the forearm where FORWARD is achieved when the wrist is fully extended. Because there is no intermediate orientation, the user can "jerk" their wrist as the orientation of the hand is constrained by the extensor tendons of the wrist. An analysis of the orientations scanned through prior to finding the target orientation found that there were a few cases where a user would deviate from the scanning sequence, e.g., they would either swap the last two orientations or would go back to a previous orientation before exploring the last two. These cases did not significantly skew the average search time, but it indicates that some users found it difficult to recall the last orientations using the cube mnemonic.

For ($o = 8$) figure 7 shows a linear increase in average search time for each consecutive orientation with the exception of the second orientation DOWN-LEFT. A Grubbs' test using consecutive differences (with the first data point transposed) revealed that this was a statistically significant outlier ($\boldsymbol{Z} = 2.00, N = 6, p < .05$). An analysis of orientations scanned through prior to finding DOWN-LEFT shows that in 44% of the cases a user would skip over this orientation and scan all the way to RIGHT-DOWN before going back to DOWN-LEFT to find this orientation. We did not find any other orientations being skipped over. An analysis of the distribution of the frequency of a user skipping RIGHT-DOWN revealed this occurred frequently to a small number of users (mean=4.05, skew=-.82, kurtosis=-.39). An explanation for some users skipping over RIGHT-DOWN is that when the user starts scanning from the start orientation (DOWN) the phone's velocity accelerates from zero to the scanning velocity which was approximated to be 84°/second. Though no false negatives were observed in preliminary trials, if the phone is accelerated too vigorously it doesn't allow for measuring its orientation until it is being rotated with a constant velocity.

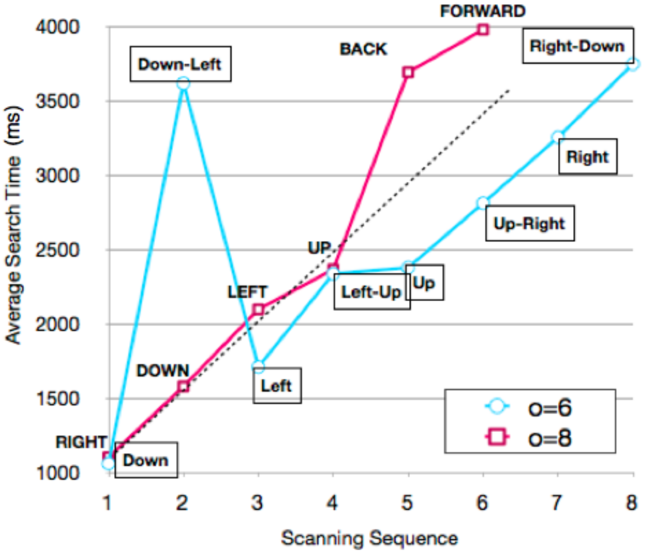

Figure 7: Comparison of mean search times for $o = 6, 8$.

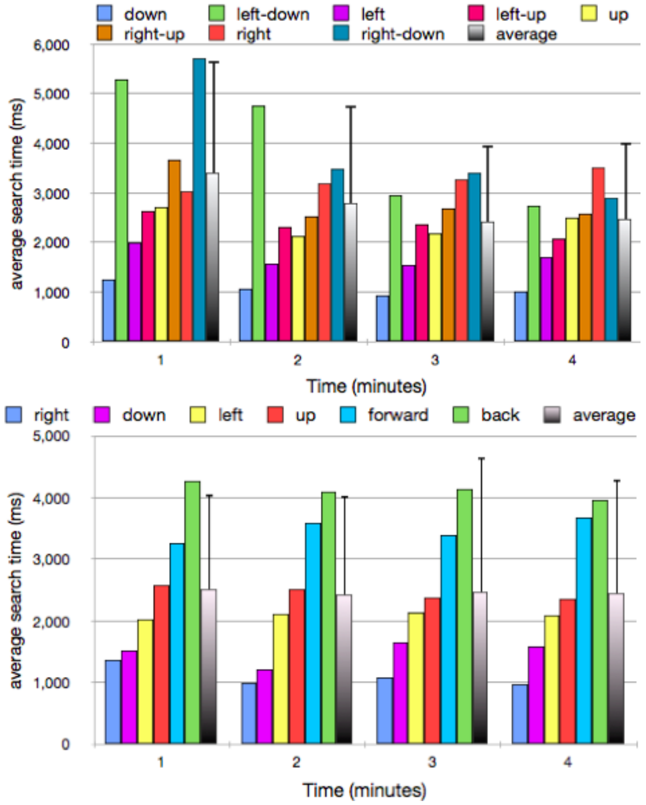

Figure 8: Average search time for each orientation per minute for $o = 6$ (top) and $o = 8$ (bottom). Last column shows the average of average search times for all orientations and its standard deviation.

Figure 8 displays the average search time for each orientation in one-minute increments for both information spaces. Though for DOWN-LEFT and RIGHT-DOWN decreases in average search times are observed over each consecutive minute, a repeated measures ANOVA with a Greenhouse-Geisser correction found these decreases not to be statistically significant ($\boldsymbol{F}_{2.016,12} = .403, p = .683$) and ($\boldsymbol{F}_{2.133,30} = 1.491, p = .248$). Using the same ANOVA no statistically significant differences in average search time were found for any orientation for any minute interval for either information space. This indicates that users becoming more familiar with the orientations for each space over time did not significantly affect search times.

To determine the performance of each information space, we average the average search time for each orientation (see Table

1 and 2) as to communicate a message any orientation is just as likely to be the target. This results in an average search time for $o = 6$ of 2,466 (SD=1,147) milliseconds and for $o = 8$ this is 2,612 (SD=934) milliseconds. For a more accurate comparison we accommodate for their differences in size, e.g., 936 milliseconds per bit for $o = 6$ and 870 milliseconds per bit for $o = 8$. This difference in average search time was not found to be statistically significant ($\boldsymbol{T}_{2,12} = .45$ $p > .05$).

5 DISCUSSION

Because proprioception is a novel output modality distinct from other modalities, we first contrast our results with previous results with analog tactile proprioceptive displays before comparing it with other forms of ear and eye free displays. Ahmaniemi [12] found that users could detect a target in a 90° horizontal range within 1.8 seconds using an average scanning velocity of 45°/second. Ahmaniemi used eight different target locations (but different target sizes) so the spatial resolutions of both proprioceptive displays are comparable. Because our orientations were defined in a significantly larger range, and higher scanning velocities were observed, the temporal resolution of discrete proprioceptive displays is presumably larger as predefined orientations are used.

Rahman [16] finds an average task time of 2 seconds for 16 levels of wrist rotations that involve pronation and supination, which uses the same range of motions as our $o = 8$. Performance is somewhat comparable though comparison is difficult as their technique is an input technique coupled to a visual display whereas our feedback is only through proprioception. Haptic feedback is most commonly used to facilitate ear and eye free interaction [7]. Chan et al [22] define seven tactons that are composed of rhythm, amplitude and frequency and which can be recognized in 2.5 seconds with a 95% accuracy. Our study did not evaluate a user's ability to correctly identify a target orientation. However, because we used the simplest of vibrotactile patterns to indicate a target orientation, identifying a hand orientation may be just as efficient or faster than recognizing a tacton as tactile feedback has a higher latency than proprioception as proprioception is an interoceptive modality [23]. Though this claim can only be evaluated using a comparative analysis, we may assume tactile and discrete proprioceptive displays offer similar performances. Because these tactons use frequency and amplitude they cannot be implemented using pager motors that are typically implemented in mobile devices. Brown and Kaaresoja [9] created nine tactons using rhythm and "simulated" amplitude by pulsing the vibrotactor where they found a recognition rate of 72%. No results are given about how fast users can identify these tactons but given their lower observed recognition rate it may be assumed that tactons that use frequency are recognized faster than those that don't. Compared with tactons, proprioceptive displays may be more efficient.

For the design of a discrete tactile-proprioceptive display we evaluated orientations in a space and in plane. Though a space allows for defining orientations that the user may be able to distinguish easier there is a small performance penalty for including orientations that cannot be scanned to with a constant velocity. To design a discrete proprioceptive display with six orientations average search time will be significantly faster when these orientations are defined in a plane. For orientations defined in a plane, we observed that some users would skip over the first orientation. An analysis of the log files did not reveal any users skipping over the first or the penultimate orientation for ($o = 6$), so either the tactile window can be increased at the cost of making it harder to distinguish consecutive orientations, but a better solution would be to distance consecutive orientations to be scanned over as such to allow for the phone to achieve a relative constant velocity rather than defining an orientation in a range where the phone is accelerating to the scanning velocity.

6 FUTURE WORK

For a complete analysis of the efficacy of proprioceptive displays, we must evaluate the user's ability to perform auto-semaphoring, i.e., whether users can successfully identify and associate different hand orientations to correspond with different messages. User studies with auto-semaphoring may significantly affect the design of the information space, for example, if orientations in a plane are found to be difficult to identify because of their proximity, we may revert to defining orientations in a space that are more distinct. On the other hand studies may show that it may be feasible to significantly increase the resolution as a related study with wrist based input [16] revealed users could distinguish 16 different orientations-though this input was coupled with visual feedback. Another concern that affects the design space is the number of different messages that can be feasibly memorized by the user. User studies were performed with the user seated and future studies will evaluate users ability to safely use a proprioceptive display in active mobile contexts, such as when driving or walking.

The performance of a discrete proprioceptive display scales up with the size of its information space. In our study we assume the type of message to be communicated is not known a priori, however, for some applications posteriori knowledge about the likelihood of a message being communicated can be collected, which can be used to optimize the efficiency of the proprioceptive display. For example, to convey the internal state change of a mobile phone, some users receive more text messages than voice mail notifications. By associating messages that are more likely to occur with orientations that are in the beginning of the scanning sequence, performance may increase significantly. Alternatively for a smaller information spaces or for sets of messages that are just as likely, users may prefer to use a random sequence or define their own sequence.

Proprioceptive displays should be considered "hybrid" displays that rely upon a small amount of haptic or audio feedback to facilitate a significantly larger information space. Potentially larger information spaces can be created when more sophisticated forms of haptic feedback are used. For example, combining a discrete proprioceptive display ($o = 8$) with the nine tactons of Brown and Kaaresoja [9] would create an information space of $\log_2(72) = 6.17$ bits. An information space of this size may be impractical to use if every state conveys a different message. Instead we could encode a parameter such as quantity using vibrotactile cues to convey analog like information such as "a lot of text messages" or "only one text message."

Most recent smart phones feature gyroscopes, which –in combination with an accelerometer– allow for determining the exact orientation of a mobile device. This opens up the opportunity to use the elongated features of the mobile device to point to a direction in a space, which may allow for users to more easily recall orientations than using the cube mnemonic that was defined for ($o = 6$).

No studies exist that evaluate the use of multiple discrete proprioceptive displays in conjunction, but using a display in each hand may have useful applications in the domain of assistive technology; for example, semaphores composed of combinations of orientations – similar to American Sign Language– could indicate different letters, which could be a useful display technique for individuals who are visually impaired. Two proprioceptive displays with six orientations each facilitates an information space that can convey an alphabet and ten numbers. A display technique such as this could be implemented using commercially available motion sensing controllers at the fraction of the cost of a tactile display. Using two proprioceptive displays in conjunction may increase the cognitive load significantly, though, users with visual impairments may be more proficient in the use of proprioceptive displays due to the plasticity effect [24]. This study may also allow for identifying whether there is a significant difference in scanning performance between the user's dominant and non dominant hand.

7 CONCLUSION

This paper explored the use of proprioception as an output modality for mobile devices through a combination of device orientation and vibrotactile feedback. As such, different hand orientations can be used to communicate messages to the user using a form of autosemaphoring. A user study with 16 users explored the temporal and spatial resolution of proprioceptive displays. The performance of a discrete proprioceptive display scales up with its size. When orientations are defined in a plane they can be scanned through with constant velocity and they can be found more efficiently than when orientations are defined in a space, on the other hand orientations defined in a space may be more easily distinguished. Orientations should not be defined in a range where the phone is accelerating to the scanning velocity as its exact orientation cannot be determined and users may skip over this orientation. The performance of discrete proprioceptive displays is comparable with more sophisticated forms of haptic feedback with the benefit that proprioceptive displays can be implemented using features currently available in most mobile devices.

8 ACKNOWLEDGEMENTS

This work is supported by NSF Grant IIS-1118074.

REFERENCES

[1] Ken Hinckley, Jeff Pierce, Mike Sinclair, and Eric Horvitz. Sensing techniques for mobile interaction. In *Proceedings of UIST 2000*, pages 91–100, New York, NY, USA, 2000. ACM.

[2] Frank Chun Yat Li, David Dearman, and Khai N. Truong. Virtual shelves: interactions with orientation aware devices. In *Proceedings of UIST 2009*, pages 125–128, New York, NY, USA, 2009. ACM.

[3] Chris Harrison and Scott E. Hudson. Minput: enabling interaction on small mobile devices with high-precision, low-cost, multipoint optical tracking. In Proceedings of *CHI '10*, pages 1661–1664, New York, NY, USA, 2010. ACM.

[4] Jungsoo Kim, Jiasheng He, Kent Lyons, and Thad Starner. The gesture watch: A wireless contact-free gesture based wrist interface. In Proceedings of *ISWC '07*, pages 1–8, Washington, DC, USA, 2007. IEEE Computer Society.

[5] Jaakko Keränen, Janne Bergman, and Jarmo Kauko. Gravity sphere: gestural audio-tactile interface for mobile music exploration. In *Proceedings of CHI 2009*, pages 1531–1534, New York, NY, USA, 2009. ACM.

[6] Scott E. Hudson, Chris Harrison, Beverly L. Harrison, and Anthony LaMarca. Whack gestures: inexact and inattentive interaction with mobile devices. In Proceedings of *TEI '10*, pages 109–112, New York, NY, USA, 2010. ACM.

[7] Ivan Poupyrev, Shigeaki Maruyama, and Jun Rekimoto. Ambient touch: designing tactile interfaces for handheld devices. In Proceedings of *UIST 2002*, pages 51–60, New York, NY, USA, 2002. ACM.

[8] Joseph Luk, Jerome Pasquero, Shannon Little, Karon MacLean, Vincent Levesque, and Vincent Hayward. A role for haptics in mobile interaction: initial design using a handheld tactile display prototype. In Proceedings of *CHI 2006*, pages 171–180, New York, NY, USA, 2006. ACM.

[9] Lorna M. Brown and Topi Kaaresoja. Feel who's talking: using tactons for mobile phone alerts. In Proceedings of *CHI 2006 extended abstracts*, pages 604–609, New York, NY, USA, 2006. ACM.

[10] E. Gardner and J. Martin. *Principles of Neural Science, fourth edition.*, chapter Coding of Sensory Information, pages 411–429. Number 21. McGraw-Hill, 2000.

[11] Bryan L Riemann and Scott M Lephart. The sensorimotor system: The role of proprioception in motor control and functional joint stability. *J Athl Train*, 37(1):80–4, Jan 2002.

[12] Teemu Tuomas Ahmaniemi and Vuokko Tuulikki Lantz. Augmented reality target finding based on tactile cues. In *Proceedings of ICMI-MLMI 2009*, pages 335–342, New York, NY, USA, 2009. ACM.

[13] Simon Robinson, Parisa Eslambolchilar, and Matt Jones. Sweep-shake: finding digital resources in physical environments. In *Proceedings of MobileHCI 2009*, pages 12:1–12:10, New York, NY, USA, 2009. ACM.

[14] Chris Harrison, Desney Tan, and Dan Morris. Skinput: appropriating the body as an input surface. In Proceedings of *CHI 2010*, pages 453–462, New York, NY, USA, 2010. ACM.

[15] Mark R. Mine, Frederick P. Brooks, Jr., and Carlo H. Sequin. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proceedings of SIGGRAPH 1997*, pages 19–26, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[16] Mahfuz Rahman, Sean Gustafson, Pourang Irani, and Sriram Subramanian. Tilt techniques: investigating the dexterity of wrist-based input. In *Proceedings of CHI 2009*, pages 1943–1952, New York, NY, USA, 2009. ACM.

[17] Tony Morelli, John Foley, and Eelke Folmer. VI-Bowling: a tactile spatial exergame for individuals with visual impairments. In *Proceedings of ASSETS 2010*, pages 179–186, New York, NY, USA, 2010. ACM.

[18] Charlotte Magnusson, Kirsten Rassmus-Gröhn, and Delphine Szymczak. Scanning angles for directional pointing. In *Proceedings of MobileHCI 2010*, pages 399–400, New York, NY, USA, 2010. ACM.

[19] Charlotte Magnusson, Miguel Molina, Kirsten Rassmus-Gröhn, and Delphine Szymczak. Pointing for non-visual orientation and navigation. In *Proceedings of NordiCHI 2010*, pages 735–738, New York, NY, USA, 2010. ACM.

[20] J M Soucie, C Wang, A Forsyth, S Funk, M Denny, K E Roach, and D. Boone. Range of motion measurements: reference values and a database for comparison studies. *Haemophilia*, Nov 2010.

[21] Andrew Crossan, John Williamson, Stephen Brewster, and Rod Murray-Smith. Wrist rotation for interaction in mobile contexts. In *Proceedings of MobileHCI 2008*, pages 435–438, New York, NY, USA, 2008. ACM.

[22] Andrew Chan, Karon MacLean, and Joanna McGrenere. Learning and identifying haptic icons under workload. In *Proceedings of the WHC 2005*, pages 432–439, Washington, DC, USA, 2005. IEEE Computer Society.

[23] K. Bark, J.W. Wheeler, S. Premakumar, and M.R. Cutkosky. Comparison of skin stretch and vibrotactile stimulation for feedback of proprioceptive information. In *symposium on HAPTICS 2008*, pages 71 –78, march 2008.

[24] L G Cohen, P Celnik, A Pascual-Leone, B Corwell, L Falz, J Dambrosia, M Honda, N Sadato, C Gerloff, M D Catalá, and M Hallett. Functional relevance of cross-modal plasticity in blind humans. *Nature*, 389(6647):180–3, Sep 1997.