

Evaluation of Tactile Feedback Methods for Wrist Rotation Guidance

Andrew A. Stanley, *Student Member, IEEE*, and Katherine J. Kuchenbecker, *Member, IEEE*

Abstract—Tactile motion guidance systems aim to direct the user’s movement toward a target pose or trajectory by delivering tactile cues through lightweight wearable actuators. This study evaluates ten forms of tactile feedback for guidance of wrist rotation to understand the traits that influence the effectiveness of such systems. We present five wearable actuators capable of tapping, dragging across, squeezing, twisting, or vibrating against the user’s wrist; each actuator can be controlled via steady or pulsing drive algorithms. Ten subjects used each form of feedback to perform three unsighted movement tasks: directional response, position targeting, and trajectory following. The results show that directional responses are fastest when direction is conveyed through the location of the tactile stimulus or steady lateral skin stretch. Feedback that clearly conveys movement direction enables subjects to reach target positions most quickly, though tactile magnitude cues (steady intensity and especially pulsing frequency) can also be used when direction is difficult to discern. Subjects closely tracked arbitrary trajectories only when both movement direction and cue magnitude were subjectively rated as very easy to discern. The best overall performance was achieved by the actuator that repeatedly taps on the subject’s wrist on the side toward which they should turn.



1 INTRODUCTION

FROM a child learning to tie her shoes to a retiree perfecting a golf swing, motor learning plays an instrumental role in almost all human activities. Mastering such skills typically requires extensive practice, which can be accelerated by guidance from a coach. Recent advances in technology have opened up the possibility of supplementing human coaching with electronic computer aids. In particular, consumer products like the Microsoft Kinect and the Nintendo Wii mote can track a user’s motions in real time with high effectiveness and low encumbrance. These technologies allow a user to practice movements in a virtual environment while a computer provides consistent feedback for an endless number of repetitions.

Beyond standard visual and audio feedback, we are beginning to see systems that deliver localized tactile cues in response to the user’s body movement. Real-time tactile feedback is typically accomplished by activating a select number of lightweight actuators worn at locations distributed across the body. In addition to traditional vibration motors, we envision tiny wearable robots that create a rich array of mechanical sensations on the skin of the human user, a new form of haptic human-robot interaction. Because tactile systems are typically small, safe, and inexpensive, we believe there is merit in exploring their potential use in human motor learning. This

article aims to provide fundamental new insights on the effectiveness of different types of tactile feedback during representative motion guidance tasks.

The next section discusses previous research on the use of wearable vibrotactile feedback to enhance motor learning, along with alternative tactile actuators for a variety of other purposes. Section 3 presents our approach for combining these two areas to create novel wearable haptic devices and haptic drive algorithms to guide wrist pronation/supination. Section 4 describes the setup of our human subject experiment and the three motion tasks it tested. The results of this study are presented in Section 5 and discussed in Section 6, and we conclude the article in Section 7. The specific contributions of this article are designs for four new tactile actuators and their respective drive algorithms, an experimental paradigm for comparing different forms of tactile feedback in motion guidance, and fundamental insights on the types of tactile cues that are most effective at guiding human motion.

2 BACKGROUND

Prior research on wearable tactile feedback includes studies on vibrotactile cues for motion guidance and work on alternative tactile actuator designs for an assortment of other applications. Here we summarize representative examples from both domains.

Within the realm of motion guidance through wearable vibrotactile actuators, both upper- and lower-limb systems have shown promising results. Bloomfield and Badler placed an array of shafted eccentric mass actuators on the user’s arm and used binary vibrations to indicate collisions while the subject tried to reach through a virtual tunnel to grasp a virtual object [1]. Focusing more explicitly on motion guidance, the Tactile Interaction for Kinesthetic Learning

- A. A. Stanley was with the Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, 19104, when this work was conducted. He is currently with Stanford University. E-mail: astan@alummi.upenn.edu
- K. J. Kuchenbecker is with the Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, 19104. E-mail: kuchenbe@seas.upenn.edu.

(TIKL) system by Lieberman and Breazeal used small voice coil actuators with graded vibration amplitude to show that vibrotactile feedback enabled significant improvements in the user’s ability to follow arm motion trajectories [2]. However, the provided feedback was not found to reduce motion tracking errors for forearm pronation/supination [2], possibly because of the saltation drive algorithm used on this joint. In contrast, Weber et al. [3] found that vibrotactile saltation feedback provided via six eccentric mass motors around the wrist outperformed verbal cues on arm rotation tasks, though task completion was slow (6 to 7 s) in both cases. Interestingly, this study showed the opposite trend for translation of the arm, with verbal feedback proving faster and more accurate than vibration, perhaps because of the low update rate of both types of cues (every 0.7 s) [3].

Our lab’s low-cost StrokeSleeve system further explores this paradigm by using shaftless eccentric mass motors to guide the user’s arm during motor rehabilitation [4]. A recent study with a new version of StrokeSleeve emphasized the importance of tuning and testing actuator drive algorithms to maintain effectiveness in tactile motion guidance [5]. For lower-limb motions, Lurie et al. proved that an array of shaftless eccentric mass motors delivering vibrotactile cues at various locations on the body can enable subjects to learn and retain modifications to their gait cycle [6]. The various actuator drive algorithms tested in this study were found to enable different levels of performance, with actuator location being easier to discern than saltation, and a pull mapping (moving toward the cue) being easier and faster to respond to than a push mapping [6].

Vibrotactile actuators have dominated tactile motion guidance thus far due to their small size, low cost, robust operation, salient cues, and simple control interface [7]; however, human skin can detect many other types of tactile cues. Expanding beyond the traditional vibrotactile approach allows a system to not only convey an alert, but also to simulate some of the physical sensations the user might feel during interactions with a real human coach. Along these lines, Baumann et al. showed that servo-motor-based mechanisms that tap and squeeze a user’s wrist can effectively emulate the attention-getting practices exhibited by humans [8]. Also using contact pressure, Tadakuma and Howe demonstrated that a whole-arm array of foam contact paddles brought into and out of contact with the arm can simulate contact with objects in virtual or remote environments [9]. In another line of work, Bark et al. created a custom device to deliver rotational skin stretch cues, aimed at helping amputees feel the location and movement of a prosthetic limb without visual attention [10]. Also focusing on skin stretch, Gleeson et al. built a portable, fingertip-mounted display that laterally deflects the distal fingerpad skin to provide directional cues in navigation

[11]. Further exploration of alternative forms of tactile feedback promises an even greater variety of future applications.

3 TACTILE FEEDBACK DESIGN

Inspired by prior work in vibrotactile motion guidance and alternative tactile device design, this study aimed to test the hypothesis that *both the tactile actuator selected and the algorithm governing its activation significantly affect the user’s performance in tactile motion guidance*. We focused on guidance of wrist pronation/supination (rotation about the forearm) because prior research has encountered mixed results in guiding this single-degree-of-freedom movement [2], [3]. By testing a large range of tactile feedback methods, we hoped to uncover additional insights into the underpinnings of tactile motion guidance and the human response to various forms and implementations of tactile cues.

This study examined ten forms of tactile feedback, spanning five tactile actuators with two drive algorithms each. The design of the devices and the selection of the parameters used in each of the drive algorithms were empirically tuned through several iterations of informal user testing, with the goal of providing clear and intuitive feedback while also being comfortable for a broad range of users. Figure 1 shows photographs of a user wearing each of the five tactile actuators used in this work. Four of them (Tapper, Dragger, Squeezer, and Twister) were custom built from 3D-printed parts and Futaba S3114 High-Torque Micro Servo motors to mimic the feel of a human tapping on, dragging across, squeezing, and twisting the subject’s wrist, as originally reported in [12]. The fifth device (Vibration) was designed to match prior research by using six shaftless eccentric mass motors (Precision Microdrives model 312-101) evenly spaced around the user’s wrist. Figure 2 shows isometric drawings of the four custom actuators, and this article’s supplemental video shows all of the actuators in use.

Each device can be controlled to move in a variety of ways; we previously used the four servo-based devices to replay recorded human movement to enable quick testing of many alternatives [12]. For this study, we focused on methods of transforming a signed wrist angle error into a graded tactile cue that guides the user to rotate their wrist to the desired orientation. As listed in Table 1, the tested drive algorithms are split into two categories: Steady, which provides a more continuous tactile cue with varying strength, and Pulsing, which delivers a repeated tactile cue with varying frequency. Larger error magnitudes are linearly mapped to higher strength cues and faster cue frequencies, saturating at 120° of error. All of the actuators except the Squeezer create bidirectional tactile signals to guide the user to rotate their wrist either

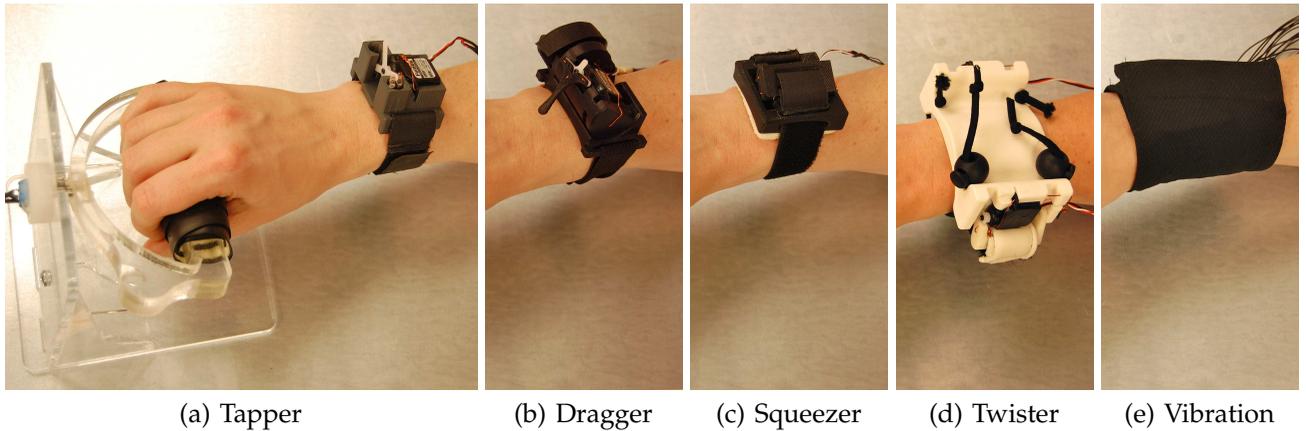


Fig. 1. The five wearable tactile devices tested in this work, along with the instrumented handle. (a) The Tapper uses a crank-slider linkage to convert servo rotation into the vertical motion of two contacting pistons spaced 30.5 mm apart. (b) The Dragger uses two servos in series to move a rubber-coated plastic tip to any position across the top of the user’s wrist or in the air above it. (c) The Squeezer uses one servo to loosen and tighten a fixed-length band around the user’s wrist. (d) The Twister mounts a pair of servos on opposite sides of the user’s wrist to move a high-friction strap back and forth on the underside of the wrist. (e) The Vibration device uses six shaftless eccentric mass motors spaced evenly around the user’s wrist inside a stretchable fabric band.

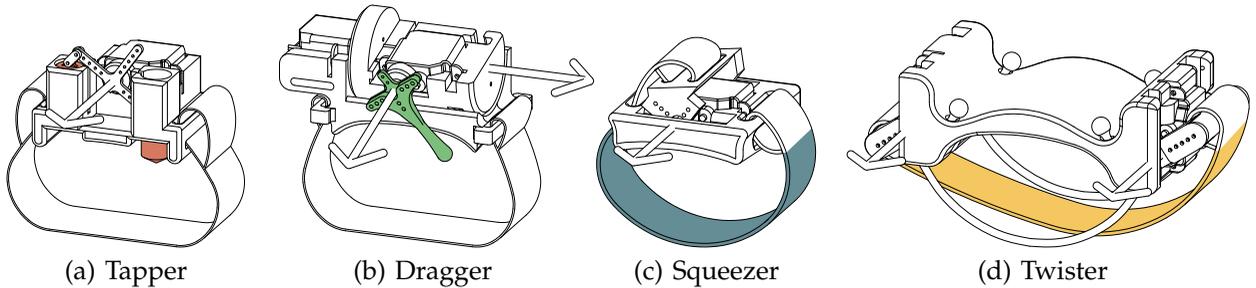


Fig. 2. Line drawings of the four novel tactile actuators. The contactor elements (Tapper pistons, Dragger tip, Squeezer band, and Twister strap) are highlighted in color, and servo rotation axes are marked by arrows.

clockwise or counterclockwise. For example, the Tapper Steady feedback presses on the top of the wrist on the side toward which the subject should rotate (pull mapping), with servo angle a linear function of wrist angle error. The chosen high-performance analog micro servos have a mass of 7.8 g, can output a maximum torque of 0.148 N·m, and can move 60° in 0.10 s when not loaded. We used an optical motion capture system in non-reported experiments to confirm that the four custom servo-based actuators closely tracked the commanded trajectories. It is important to note that none of these actuators creates a physical torque that can rotate the user’s hand toward the target angle. Instead, the tactile feedback works through a cognitive channel, so we can be confident that all observed wrist movement is due to the user’s action.

As seen in Fig. 1(a), a plastic handle with a built-in rotational potentiometer allows real-time measurement of the user’s wrist angle with a resolution of 0.3° . 50 times per second, a MATLAB script running on a PC measures the wrist potentiometer angle using a Phidget 8/8/8 USB Interface Kit; this angle is compared with the target wrist angle to calculate the angular error, which has both magnitude and direc-

tion. For the Tapper, Dragger, Squeezer, and Twister, the script then uses a Phidget Advanced Servo 8-Motor USB board to send pulse-width modulation (PWM) signals to reposition the actuator’s servo(s). For the Vibration device, it uses PWM on the Phidget 8/8/8 to adjust the six motor voltages. Throughout this study, the tactile feedback was programmed to turn off completely when the user was within 7.5° of the target; this 15° deadband width corresponds well with the $0.64 \pm 6.5^\circ$ average error in forearm pronation/supination positioning for healthy controls using their dominant arm without visual feedback [13]. Smaller deadbands were empirically tested, but users often became confused when slight movements back and forth across the deadband resulted in the feedback quickly turning on and off in opposite directions.

4 EXPERIMENTAL METHODS

We conducted a human subject experiment to test the ten forms of tactile wrist-motion guidance described in the previous section. Each subject used each form of feedback to complete three tactile guidance tasks, as outlined in Table 2 and detailed below. The three tasks were presented in the same order for each form

Actuator	Steady Algorithm	Steady Range	Pulsing Algorithm	Pulsing Range
Tapper	Servo rotates to press down on the side toward which the user should turn.	30° – 55°	Servo follows a 35° repeating half-sinusoid to tap on the side toward which the user should turn.	2 Hz – 6 Hz
Dragger	Tip drags smoothly across the user’s wrist in the direction the user should turn, resetting quickly.	1 Hz – 4 Hz	Tip taps six times across the user’s wrist in the direction the user should turn, resetting quickly.	3 Hz – 12 Hz*
Squeezer	Servo rotates to tighten the strap around the wrist. (non-directional)	10° – 45°	Servo follows a 25° repeating half-sinusoid, squeezing tightest at the peaks. (non-directional)	2 Hz – 6 Hz
Twister	Both servos rotate to stretch the wrist skin in the direction in which the user should turn.	15° – 40°	Both servos repeatedly stretch the skin in the direction in which the user should turn, resetting 3x more slowly.	2 Hz – 4 Hz
Vibration	Motor vibrates on top of wrist on the side toward which the user should turn, graded by PWM.	45% – 100%	All six eccentric mass motors vibrate sequentially in the direction in which the user should turn.	3 Hz – 7.8 Hz*

TABLE 1

The ten drive algorithms tested in this work, including one Steady and one Pulsing for each device. All except the Squeezer deliver directional tactile cues that depend on the sign of the wrist angle error. The lower value in each range indicates the device’s output at the edge of the $\pm 7.5^\circ$ deadband, while the upper value is the output at or above 120° of error; the cue scales linearly with error between these extremes. Inside the deadband, the algorithms all deliver no tactile cues (0° , 0 Hz, or 0%). *Note that Dragger Pulsing and Vibration Pulsing specify the number of tactile events (taps or buzzes) per second, not the number of full cycles per second.

Task	Objective	Sets	Metrics
Directional Response	Wait for the cue to start, then move accurately yet quickly in the direction it indicates.	12 trials (Practice) 12 trials (Test)	Reaction time and directional error rate
Angle Targeting	Rotate your wrist to the target angle, guided only by tactile cues, then hold still for 1 second.	8 trials (Practice) 8 trials (Test)	Rise time, settling time and maximum overshoot
Trajectory Following	Rotate your wrist to track a continuously varying target angle, guided only by tactile cues.	30 seconds (Test 1) 30 seconds (Test 2)	RMS error

TABLE 2

The three tactile motion guidance tasks tested in the study. The tasks progress from easy to difficult and were always conducted in this order, grouped by tactile feedback type. Each subject did twenty total sets of each task.

of feedback, increasing in complexity from (1) directional response, to (2) angle targeting, to (3) trajectory following. The constant task order within each form of feedback helped the subject learn to interpret cue direction and magnitude before undertaking the challenging trajectory following task. The video associated with this article shows three directional response trials for each form of tactile feedback, plus one example of angle targeting and one of trajectory following. Through the diversity of the requirements in each task, the study aspired not only to quantify the usefulness of each feedback option for these specific tasks, but also to elucidate characteristics of each form of tactile feedback that could guide actuator and algorithm selection for a variety of tactile motion guidance applications.

This within-subjects study enrolled $N = 10$ individuals (9 right handed and 1 left handed, 7 male and 3 female, age range from 20 to 30 years, average age of 22.2 years) drawn from the general university population. It took each subject about two hours to complete two sets of each task for all ten forms of feedback. The ten feedback types (5 devices \times 2 drive algorithms) were presented in a pseudorandom order to eliminate presentation order bias. The ordering fit a ten-by-ten randomized Latin square, so that among the ten subjects, exactly one subject tried each form of



Fig. 3. The experimental setup. The subject wears a wrist-mounted tactile actuator and holds the instrumented handle while closing their eyes and listening to white noise. Key presses are entered on the computer keyboard with the other hand (not visible).

feedback first, last, and at every position in between, with a random set of transitions between feedback forms. Subjects were permitted to take a break at any time between sets to minimize mental fatigue.

Only subjects over the age of 18 with normal sensorimotor abilities could participate in this study. All subjects gave informed consent and were compensated with a \$25 gift card. Subjects completed a demographics questionnaire at the start of the experiment.

The experimenter helped the subject don each device on the wrist of their dominant hand and adjusted it to fit correctly. The subject then used their dominant hand to hold the instrumented pronation/supination handle, which was fixed to the desk with adhesive tape (Fig. 3). Before starting the three tasks, the subject pronated and supinated his or her wrist to its maximum comfortable range of motion to calibrate the range of target angles. Subjects wore a set of over-the-ear headphones playing white noise to mask out auditory cues from the devices and any distractions from the lab environment, and they kept their eyes closed during testing to eliminate any visual cues from the devices. There was a survey after each form of tactile feedback, and one final survey at the end for a total of twelve questionnaires. The University of Pennsylvania Institutional Review Board approved all testing procedures under Protocol #814128.

4.1 Directional Response Task

This task was designed to measure the clarity with which each type of tactile feedback communicates the desired direction of wrist rotation. Each trial started with the subject holding his or her wrist in the center of its rotational range of motion, typically with the thumb up and the handle approximately vertical. The subject indicated he or she had centered his or her wrist by pressing a key on the keyboard with the non-dominant hand. After a random delay between 2 and 5 seconds, the tactile actuator would begin delivering a cue in one of the two directions (clockwise or counterclockwise). The magnitude of the cue remained constant across all trials at 75% of the maximum strength or frequency; this value was chosen through pilot testing to ensure it was easily detectable without being startling. Subjects were instructed to quickly turn their wrist in the direction that the tactile cue indicated, prioritizing directional accuracy over speed. After the subject had turned 45° from the starting position in either direction, the cue turned off and the headphones beeped to signal the end of the trial. The subject then recentered his or her wrist and pressed a key when ready for the next trial. Subjects were allowed to correct the direction of their motion if they realized they had initially responded incorrectly but had not yet passed the 45° threshold. Because the Squeezer cannot create a bidirectional signal, subjects were asked to select a direction at random for this device. We included these trials in the study to give subjects equal exposure to all ten types of feedback.

The twelve trials in each directional response set contained six cues in each direction in random order. Subjects were informed that the directions would be random, but not that they would be balanced across sets. The first response task set was treated as a practice set and was not included in the data analysis. At the start of this set, subjects were allowed to open their eyes to watch the device, remove the headphones, and

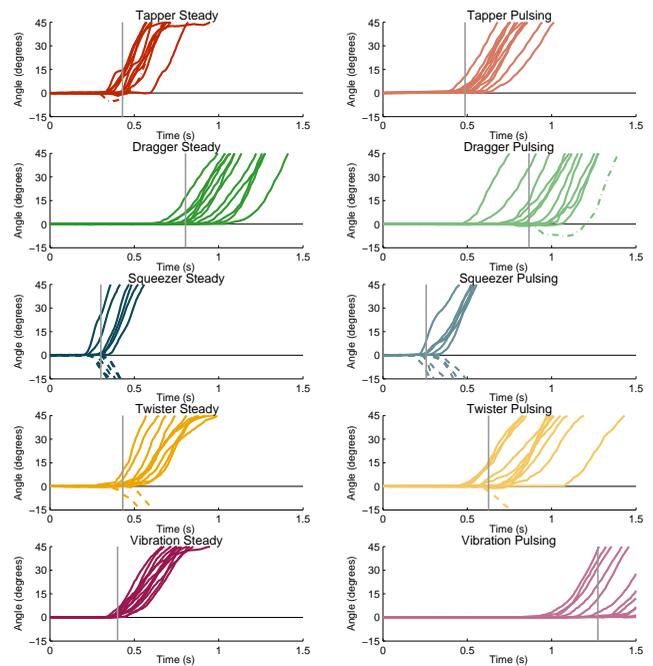


Fig. 4. Sample normalized data for the directional response task with all ten types of tactile feedback. The cue begins at time zero. Solid lines are correct responses, dashed lines are incorrect, and dash-dotted lines are trials where the subject corrected the direction of movement. The vertical gray lines show the median reaction time for each tactile feedback type.

even ask the experimenter clarifying questions to gain intuition for how to respond to each new tactile cue. Subjects were encouraged to close their eyes and put the headphones back on by the end of the practice set. The second set constituted the test set and was conducted without any visual or auditory input.

Figure 4 shows sample directional response data for one subject with all ten types of tactile feedback; all motions are plotted such that movement in the direction of the cue is positive. We calculated reaction time as the time elapsed between the initiation of the tactile cue and the user's first deviation from the waiting position by more than 1° regardless of whether the motion later changed direction; visual inspection of the data indicated that 1° sufficed to prevent any false reaction time readings from unintentional movement. Furthermore, all recorded reaction times were found to be larger than 0.17 s, so none were discarded as anticipatory false positives [14]. For each subject we calculated the median reaction time from the twelve trials in the set, the interquartile range (IQR) of these twelve reaction times, the proportion of trials in which the subject initially moved in the correct direction, and the proportion of trials in which the subject crossed the 45° threshold on the correct side.

4.2 Angle-Targeting Task

After completing two sets of the directional response task, each subject began the angle-targeting task for

the same device and drive algorithm. This task was designed to test the effectiveness of each type of tactile feedback at guiding the subject to move to a steady wrist angle. The subject centered his or her wrist and started the first trial with a key strike. The computer then selected a random target angle that was within the subject's range of motion and also at least twenty percent of the range of motion away from his or her current wrist angle. Based on the direction and grading of tactile cues described in Section 3, the subject then had to rotate his or her wrist to the target angle and stay within the tolerance of the deadband for at least one second. After one second in the deadband with no tactile cues, the headphones would beep to indicate the completion of the trial, and the computer would select a new target angle, commencing the next trial immediately. The headphones beeped multiple times to indicate the completion of a set (after eight trials). Like the directional response task, the position targeting task consisted of one practice set, to familiarize the subject with responding to both the magnitude and the direction of the tactile cue, and one test set. Figure 5 shows the series of target angles and the angular trajectory of the subject's wrist for one sample set with one form of tactile feedback.

To better analyze the data from the angle targeting task, we viewed each individual movement to a target position as a response to a step input. Subtracting the starting angle and normalizing each trial by the angular distance and direction between targets converts each trial into a unit step response. Figure 6 plots the test data from one sample user in this normalized step-response form. Inspired by classical control theory, we measured the rise time, maximum overshoot, and settling time for each trial and calculated the median of each of these values for each user for each form of tactile feedback. We measure rise time as the number of seconds it took the subject to move from ten percent of the distance to the target to ninety percent of the distance to the target; if the deadband covers more than 10% of the distance to the command, we measure rise time to the edge of the deadband instead of the 90% point, since that is where the tactile

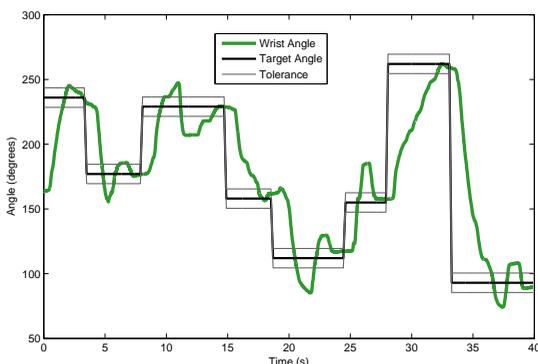


Fig. 5. Sample raw data for the angle-targeting task with one type of tactile feedback.

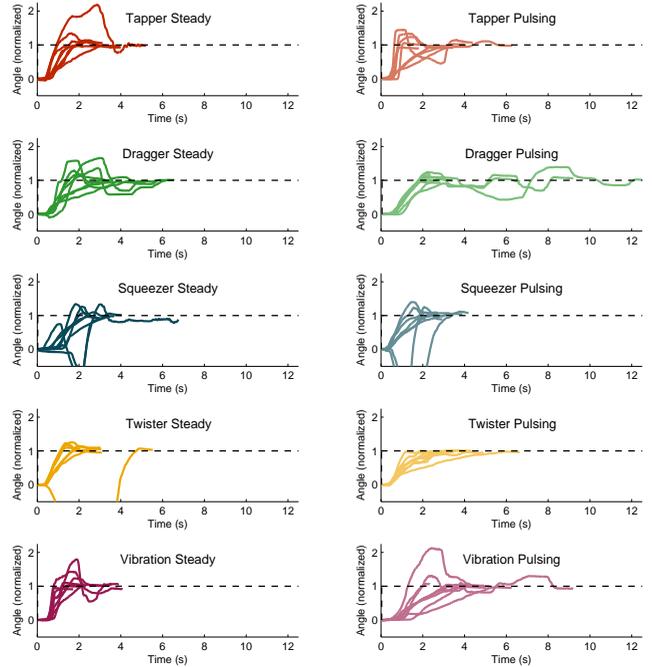


Fig. 6. Normalized sample data for the angle-targeting task with all ten types of tactile feedback.

feedback turns off. The furthest distance turned past the target angle, measured as a proportion of the step magnitude, defines the max overshoot. The settling time equals the number of seconds required for the user to enter the deadband and stay there, which is also equal to the length of the trial minus one second.

4.3 Trajectory-Following Task

Last, each subject completed two sets of a trajectory-following task with each form of tactile feedback. This task was designed to determine how well the chosen actuator and drive algorithm could guide the subject in following an arbitrary wrist angle trajectory. Having already experienced the tactile signals from the device, the subject was not permitted a practice set for this task. Instead, both sets were included in the data analysis. Each set consisted of a thirty-second-long continuous trajectory of target angles, starting at the center of the subject's range of motion. An algorithm generated a random combination of sine waves of varying magnitudes for the trajectory, with constraints to keep the standard deviation of the angles in each trajectory between 31 and 38° , and the average angular speed of each trajectory between 9 and 11° per second. These empirically chosen limits were enforced to keep the difficulty of the trajectory approximately constant while always providing new trajectories.

Figure 7 shows a sample set of trajectories and the resulting wrist angle movement from one subject for each form of feedback. For each trial, we measure the root-mean-square (RMS) error between the user's trajectory and the edge of the deadband, such that angles within the deadband contribute no error. RMS error punishes large deviations from the target trajectory.

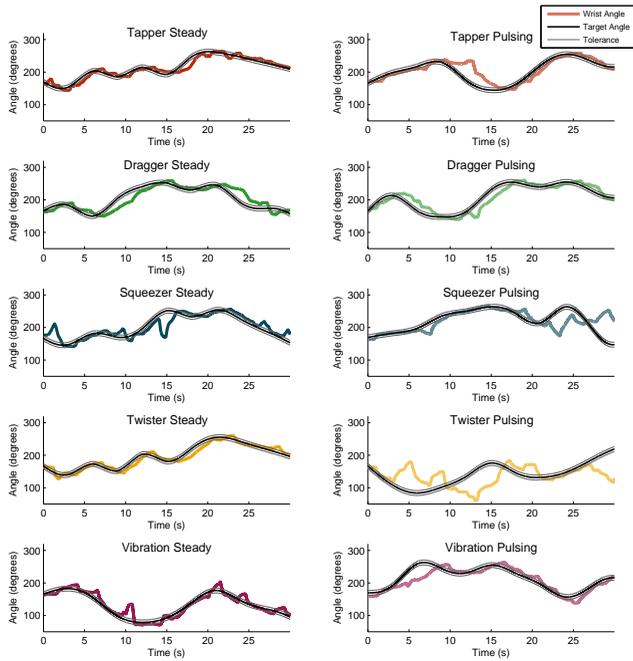


Fig. 7. Sample data for the trajectory following task with all ten types of tactile feedback.

5 RESULTS

We analyzed the data for most collected metrics using three-way analysis of variance (ANOVA) with first-order interactions and post-hoc pairwise comparisons. The three ANOVA factors were device type (Tapper, Dragger, Squeezer, Twister, or Vibration), algorithm type (Steady or Pulsing), and subject number. Device and algorithm were treated as fixed effects, while subject was treated as a random effect. Following Whelan’s advice for analyzing reaction time data [14], we performed the ANOVA on the logarithm of the reaction time, rise time, and settling time data to normalize their variance. The same transformation was applied to the trajectory-following RMS error, as its variance was also found to scale with magnitude. Lilliefors tests on the transformed data confirmed the normality assumption. We discuss the normality of other types of data in the detailed results sections below. In all cases, the box plots show the untransformed data for ease of interpretation.

Our first null hypothesis, H_{DS+AS} , states that *neither of the fixed-random interaction variance components is larger than zero*. If this test shows no significance, as found for all of the metrics from our experiment ($\alpha = 0.05$, values not reported), we do not reject H_{DS+AS} , and the statistical analysis simplifies greatly.

Our second null hypothesis H_{10} states that *the ten treatment combinations are all the same*. If the p-value for H_{10} falls below α , we reject the hypothesis that the ten feedback conditions are the same, and we continue further with the analysis to test the more specific null hypothesis H_{DA} that *the interaction between the two fixed effects (device and algorithm) is zero*. If the interaction

between device and algorithm is significant, we use the method described below to compare each of the ten forms of tactile feedback against one another, essentially reducing the test to a two-way ANOVA with one fixed variable (device combined with algorithm) and one random variable (subject).

After testing the interaction between device and algorithm, we test two final null hypotheses about the main effects: H_D is *the five devices do not differ from one another*, and H_A is *the two algorithm types do not differ from each other*. If we reject the first of these hypotheses, we do post-hoc tests to compare the five devices against one another. If the second is rejected, we know that the two algorithms achieved significantly different means, and we use the comparison test only to verify the order.

For all metrics where we conclude that the ten feedback conditions are not all the same, we run a variety of multiple comparison tests based on the results of the final three hypotheses. We conclude that two treatments or treatment combinations are different if and only if the absolute difference between the corresponding means exceeds a critical value (cv) times the standard error of the difference between the means. We use Tukey’s method and divide by $\sqrt{2}$ to calculate cv for each test.

Table 3 gives the value of the F statistic, its degrees of freedom, and the associated p-value for all of the tests we ran on the metrics analyzed in the study. We use $\alpha = 0.05$ to determine significance, and we report the size of all effects using the η^2 method. The following subsections summarize the results for the three tasks and the questionnaire, including box plots of the data and post-hoc tests when appropriate. All box plots group the tactile actuators by similar colors, with the lighter color indicating the Pulsing drive algorithm. The box marks the interquartile range (IQR), the center line marks the median, the whiskers mark the range up to 1.5 times the IQR, and x’s mark outliers (data points that fall beyond the whiskers). Extreme outliers are clipped down to the dotted line in some cases to preserve the visibility of the plot.

5.1 Directional Response

The directional response task shows the speed and accuracy with which subjects responded to the ten different tactile cues. Figure 8 shows box plots for the metrics measured in this task: reaction time median, reaction time IQR, and directional error rate. Though subjects were allowed to correct their initial response direction, this rarely happened. Thus the incorrect initial response box plot (not shown) is almost identical to the incorrect final response box plot shown.

Median of reaction times As listed in Table 3, the median of the transformed reaction time data reached significance on all four null hypotheses, H_{10} , H_{DA} , H_D , and H_A . Multiple pairwise comparison tests on the ten forms of feedback showed a clean split into three

Task	Metric	H_{10}	H_{DA}	H_D	H_A
		Ten Feedback Types Do Not Differ	Device \times Algorithm Do Not Interact	Five Devices Do Not Differ	Two Algorithms Do Not Differ
Directional Response	Median of Reaction Times	F(9,81) = 33.16 p < 0.0001, $\eta^2 = 0.73$	F(4,81) = 7.69 p < 0.0001, $\eta^2 = 0.08$	F(4,81) = 60.07 p < 0.0001, $\eta^2 = 0.59$	F(1,81) = 27.33 p < 0.0001, $\eta^2 = 0.07$
	IQR of Reaction Times	F(9,81) = 13.64 p < 0.0001, $\eta^2 = 0.54$	F(4,81) = 4.06 p = 0.0048, $\eta^2 = 0.07$	F(4,81) = 20.82 p < 0.0001, $\eta^2 = 0.36$	F(1,81) = 23.29 p < 0.0001, $\eta^2 = 0.10$
Angle Targeting	Median of Rise Times	F(9,81) = 3.27 p = 0.0019, $\eta^2 = 0.17$	F(4,81) = 2.79 p = 0.0317, $\eta^2 = 0.07$	F(4,81) = 1.64 p = 0.1710, $\eta^2 = 0.04$	F(1,81) = 11.73 p = 0.0010, $\eta^2 = 0.07$
	Median of Settling Times	F(9,81) = 4.65 p = 0.0001, $\eta^2 = 0.32$	F(4,81) = 3.28 p = 0.0152, $\eta^2 = 0.10$	F(4,81) = 2.59 p = 0.0427, $\eta^2 = 0.08$	F(1,81) = 18.36 p = 0.0001, $\eta^2 = 0.14$
Trajectory Following	RMS Error	F(9,81) = 5.51 p < 0.0001, $\eta^2 = 0.34$	F(4,81) = 9.07 p < 0.0001, $\eta^2 = 0.25$	F(4,81) = 3.27 p = 0.0155, $\eta^2 = 0.09$	F(1,81) = 0.28 p = 0.5980, $\eta^2 = 0.00$
	Intuitiveness	F(9,81) = 4.57 p = 0.0001, $\eta^2 = 0.26$	F(4,81) = 0.81 p = 0.5223, $\eta^2 = 0.02$	F(4,81) = 9.11 p < 0.0001, $\eta^2 = 0.23$	F(1,81) = 1.45 p = 0.2319, $\eta^2 = 0.01$
Questionnaire	Discerning Direction	F(9,81) = 6.68 p < 0.0001, $\eta^2 = 0.37$	F(4,81) = 2.23 p = 0.0725, $\eta^2 = 0.05$	F(4,81) = 12.59 p < 0.0001, $\eta^2 = 0.31$	F(1,81) = 0.86 p = 0.3554, $\eta^2 = 0.01$
	Discerning Strength	F(9,81) = 6.02 p < 0.0001, $\eta^2 = 0.34$	F(4,81) = 4.85 p = 0.0015, $\eta^2 = 0.12$	F(4,81) = 6.88 p = 0.0001, $\eta^2 = 0.17$	F(1,81) = 7.25 p = 0.0086, $\eta^2 = 0.05$

TABLE 3

Results of three-way ANOVA on the primary metrics measured in the study. The null hypotheses are listed in the column headings. Bold cells contain statistically significant results, using $\alpha = 0.05$, indicating that we can reject the corresponding null hypothesis. Appropriate post-hoc comparisons are reported in the text.

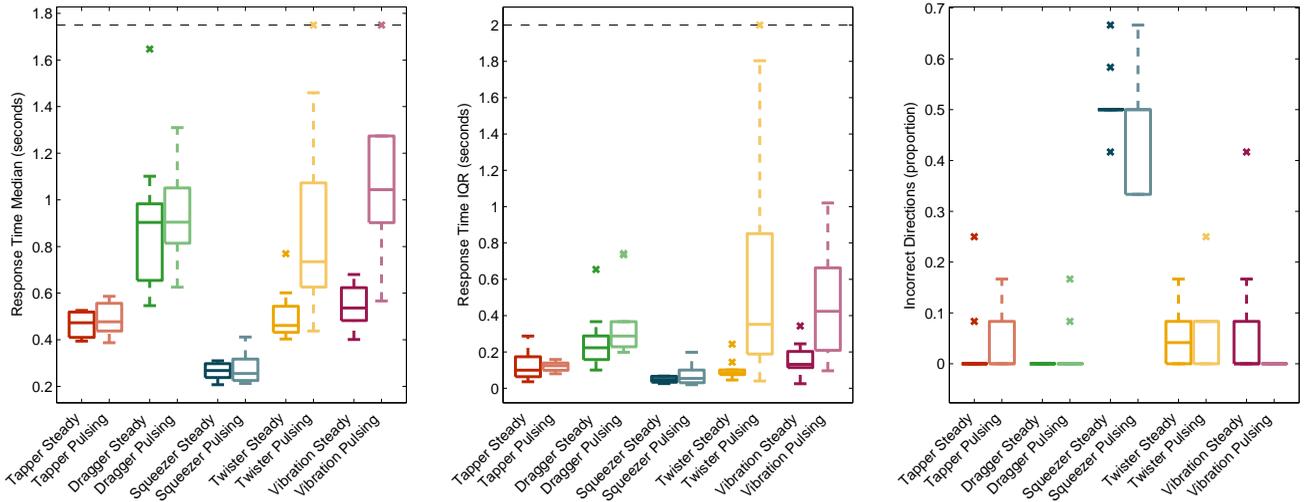


Fig. 8. Reaction time median, reaction time IQR, and error rate results for the directional response task.

separate groups, with Squeezer Pulsing and Squeezer Steady in the group with fastest reaction times, and Dragger Steady, Dragger Pulsing, Twister Pulsing, and Vibration Pulsing in the group with the slowest reaction times. Comparing the five devices shows the Squeezer had the shortest reaction times, followed by the Tapper. Of the remaining three devices, the Twister had significantly faster reaction times than the Dragger, but Vibration did not show significant differences from either of those two. Comparing the two algorithm types shows reaction times were 0.27 s faster on average for Steady algorithms.

IQR of reaction times The IQR of the transformed reaction time data reached significance on all four null hypotheses, H_{10} , H_{DA} , H_D , and H_A . Multiple pairwise comparison tests on the ten forms of feed-

back showed that Twister Pulsing had significantly larger IQRs than all other devices except Vibration Pulsing and Dragger Pulsing; otherwise there were no significant differences. Among the devices, the Twister was found to have larger IQRs than the Tapper and the Squeezer. Among algorithms, Pulsing had IQRs on average 0.22 s higher than Steady.

Directional error rate Lilliefors tests on the error rate data showed that this measurement did not satisfy the ANOVA assumption of normality, evidently because it included a very large number of zeros. We were not able to find a transformation that enabled it to pass the Lilliefors tests, so we opted not to report any formal statistical analysis on directional error rate. Instead, we invite the reader to visually examine the right-most box plot in Fig. 8. The two Squeezer

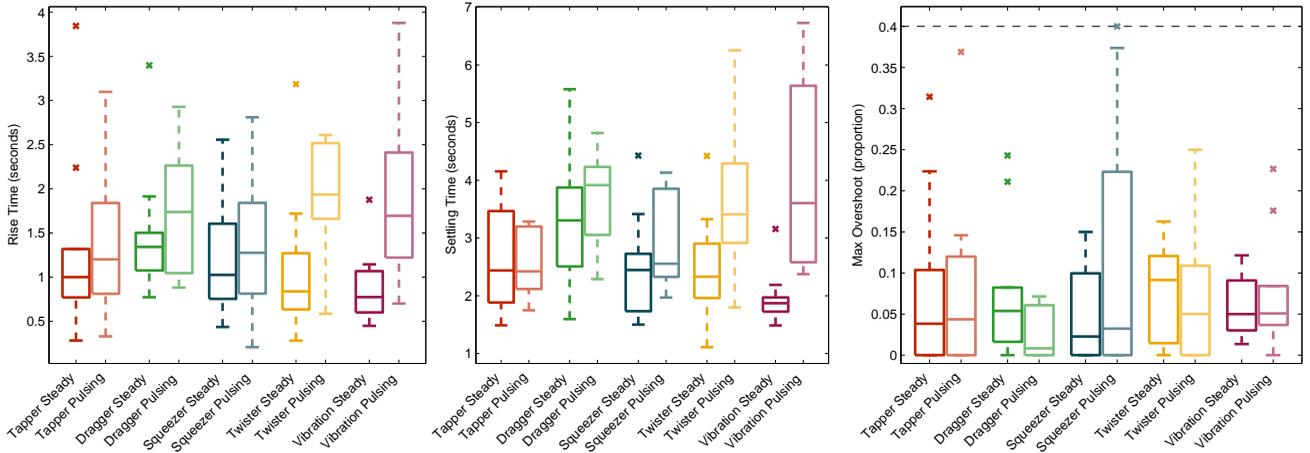


Fig. 9. Rise time, settling time, and maximum overshoot results for the angle-targeting task.

distributions both have medians of 0.5, while the other eight distributions all have medians of 0.00 or 0.04. We conclude that the Squeezer had a higher error rate than the other four devices, with no other obvious differences between the tested conditions.

5.2 Angle Targeting

The angle-targeting task tests the subject’s ability to move to target wrist angles guided only by tactile cues. Figure 9 shows box plots for median rise time, settling time, and maximum overshoot for each subject’s performance on this task.

Rise time Analysis of the transformed rise time data showed significant differences between the ten feedback forms, significant device-algorithm interactions, no significant differences between the five devices, and significant differences between the two algorithms. A multiple comparison test between the ten treatment combinations showed significantly faster rise times for Vibration Steady than for both Vibration Pulsing and Twister Pulsing, but no other statistically distinguishable differences between the remaining feedback forms. Comparison of the algorithms showed Steady is on average 0.47 s faster than Pulsing.

Settling time The transformed median settling time data from each subject showed statistical significance for all four of the null hypotheses. Comparing all ten combinations showed subjects settled on a target angle faster for Vibration Steady than for Dragger Steady, Dragger Pulsing, Twister Pulsing, and Vibration Pulsing. Twister Steady also achieved significantly faster settling times than Vibration Pulsing. No other pairwise differences were significant. Comparing the five devices against one another showed that on average the subjects settled on the target angle significantly faster with the Tapper than with the Dragger; otherwise no significant differences were found between devices. The Steady algorithms resulted in settling times 0.79 s faster on average than Pulsing feedback.

Maximum overshoot As with directional error rate, the maximum overshoot measurements contain

many zeros and were found not to obey a normal distribution, so we did not conduct an ANOVA on this data set. Instead, we visually examine the right-most panel of Fig. 9 and observe that there are no apparent differences between the ten forms of feedback for median maximum overshoot in angle targeting.

5.3 Trajectory Following

The trajectory-following task sought to elucidate the overall effectiveness of each form of tactile feedback in guiding arbitrary wrist movements. We averaged the two RMS errors from each subject’s trajectory-following trials before running this ANOVA. Figure 10 displays the box plot for this task metric.

RMS error There were significant differences between the ten forms of feedback, significant device-algorithm interactions, significant differences between the five devices, and no significant differences between the two algorithms for transformed RMS error. Comparing all ten combinations showed significantly lower RMS errors for Tapper Pulsing and Dragger Steady than for Tapper Steady, Squeezer Pulsing, and Twister Steady. RMS error for Vibration Steady was also lower than Twister Steady. Otherwise no distinctions between the remaining forms of

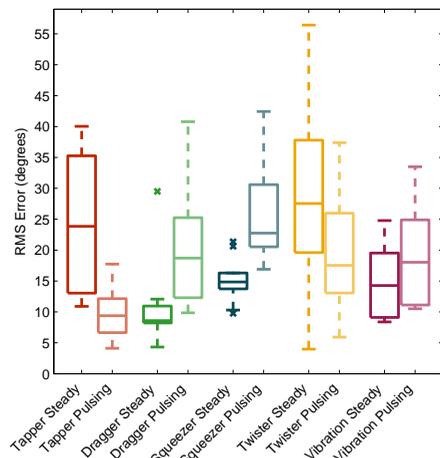


Fig. 10. RMS error results for trajectory following.

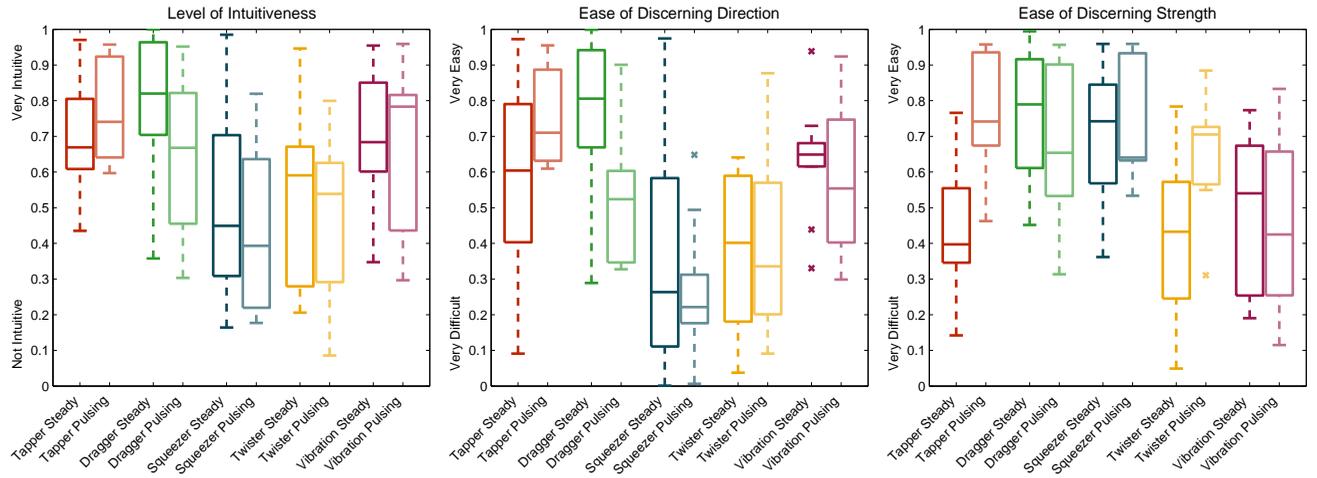


Fig. 11. Subjective ratings of tactile feedback intuitiveness, ease of discerning tactile cue direction, and ease of discerning tactile cue strength from the questionnaires for every combination of actuator and drive algorithm.

feedback could be made. The only significant difference between devices was that the Dragger had lower RMS errors than the Twister.

5.4 Questionnaire

After completing all three tasks with a given form of tactile feedback, the subject answered eight subjective questions on a paper survey. Ratings were indicated by placing a vertical slash along a horizontal line with labeled endpoints; unlike Likert-style responses, this visual analog scale yields interval data that can be analyzed via ANOVA [15]. Figure 11 shows box plots of the responses to the three key questions.

Intuitiveness When subjects rated the level of intuitiveness of the tactile feedback, significant differences were found between the ten feedback forms and between the five devices, but not for the device-algorithm interaction or the two algorithms. Comparing the five devices showed that the Tapper, the Dragger, and the Vibration actuator were significantly more intuitive than the Twister or the Squeezer.

Discerning Cue Direction The same hypotheses achieved significance when subjects rated the ease of discerning the direction of the feedback cue. As with intuitiveness, the Tapper, Dragger, and Vibration created direction signals that were easier to discern than those of the Twister or Squeezer.

Discerning Cue Strength Ratings of the ease of discerning the strength of the signal showed significance on all four of the null hypotheses, H_{10} , H_{DA} , H_D , and H_A . Subjects found the signal strengths easier to discern from the Dragger and Squeezer than from the Twister or Vibration, and they also felt that Pulsing algorithms created more easily distinguishable signal strength than Steady algorithms.

6 DISCUSSION

This experiment yielded a wealth of quantitative and qualitative data about the ten tested forms of tactile motion guidance. Aided by the experimenter's

observations and the subjects' written and verbal comments, here we aim to explain and interpret the trends uncovered in the data.

6.1 Directional Response

In the directional response task, both Squeezer algorithms resulted in reaction times significantly shorter than those from any other forms of feedback. The Squeezer is the only device that does not require the user to discern the direction of the cue; as expected, its reaction times closely match the minimum human reaction time to external stimuli [14], indicating that our testing system did not include any significant delays. The high directional error rates for the Squeezer also come as no surprise, given the 50% chance of guessing correctly. The other four devices all achieved error rates close to zero during the test sets, indicating that subjects were easily able to discern cue direction. Our observations during informal trials and this study suggest that it takes subjects only a handful of practice trials to reach this high level of directional accuracy.

As seen in Fig. 8, the significantly faster reaction times for Steady algorithms seem to stem primarily from the Twister and Vibration. Whereas subjects primarily reacted to the initial cue from the Steady algorithms, they often waited through multiple Pulsing cycles before deciding which direction to move. Because the onset of each tactile cue represents the most salient portion of the feedback, Steady cues offered little incentive to wait beyond the initial stimulus. Some forms of feedback, like Vibration Pulsing and both of the Dragger algorithms, convey direction via tactile cues that move across the skin; these designs forced users to wait beyond the cue onset to discern the direction, resulting in longer reaction times. While both Tapper Pulsing and Twister Pulsing convey cue direction immediately, only Tapper Pulsing resulted in reaction times matching those of its Steady counterpart. Twister Pulsing lagged significantly behind Twister Steady, possibly because the sensation of the

band returning to center confused users or dulled the directional sensation. The IQR of reaction time shows how repeatable subjects were in reacting to the different tactile stimuli. The IQR results largely mirror those for the median of reaction time; the Twister Pulsing feedback created the largest mean dispersion of reaction times, while the Tapper and Squeezer devices were found to have the smallest mean dispersion. We thus conclude that tactile feedback systems should use the location of skin contact, the location of vibration, or non-pulsing lateral skin stretch to convey movement direction when the speed of response is most important.

6.2 Angle Targeting

The angle-targeting task was more complex than the directional response task because both the direction and the magnitude of the tactile feedback cue varied. From our qualitative observations, subjects seemed to focus on one of these two indicators to accomplish the task, but not necessarily the same one for each form of feedback. The speed of a subject's rise time often seemed to depend on their confidence in interpreting that specific tactile cue. Slower rise times for feedback forms like Twister Pulsing and Vibration Pulsing tend to match the slower reaction times those same forms of feedback achieved in the directional response task. If the subject needed a long time to decide which direction to move, they were likely not confident enough in their interpretation to move quickly in the angle-targeting task, so they relied more on the grading of the cue as they approached or moved away from the target angle. We conclude that cues that are directionally easy to interpret seem to enable subjects to move more quickly to an unknown target angle.

Digging deeper into this data, we see that the rise times for Steady algorithms were significantly faster than those for Pulsing algorithms. We believe this difference stems from the fact that the deadband around the target is much more salient during Steady algorithms than during Pulsing algorithms. Most subjects probably learned during the practice set that moving quickly through the deadband with a Pulsing algorithm feels much like a pulse, since the feedback turns off and then on in the opposite direction. Such an observation would lead subjects to slow their movements, as seen in the data. With a Steady drive algorithm, a momentary pause in the feedback could indicate only that the user had reached the target tolerance, leading to faster rise times. This subtlety may also help explain the lack of any obvious differences in the maximum overshoot data. One would expect faster rise times to lead to larger overshoots, but the salience of the deadband helps the subject quickly halt their movement, giving overshoots similar to those that occur with slower rise times. These trends can be summarized by saying that an easy-to-discern deadband allows the subject

to move toward a target more quickly, and that our implementation of the deadband was easiest to feel during Steady algorithms.

We observed that long settling times typically occurred when a subject started to feel lost and began questioning his or her initial interpretation of the cue. This situation often caused the subject to move his or her wrist back and forth repeatedly to elicit the most salient tactile signals available (strong directional cues and/or the deadband). As with rise time, the extra salience of the deadband during Steady drive algorithms helped reduce settling times, particularly with the Vibration Steady feedback, in which the activation and deactivation of the motors seemed particularly noticeable. The Tapper was also found to outperform the Dragger on settling time. These findings reiterate the importance of choosing tactile motion guidance cues that clearly convey the intended movement direction along with close proximity to the target.

The biggest advantage of the Pulsing algorithms seemed to be that subjects were less likely to become completely lost. With a Steady algorithm, a subject who misses or fails to interpret the initial cue has to pass through the deadband to feel the signal restart. On the other hand, a Pulsing algorithm cycles through its signal regularly without forcing the user to move, providing a level of redundancy that may be valuable in certain applications. Furthermore, the questionnaires showed that the varying frequencies of the Pulsing algorithms were more distinguishable than the varying strengths of the Steady algorithms. Thus, when a subject struggled to interpret the direction of a certain cue, they could rely on the speeding or slowing of the Pulsing signal to indicate whether they were moving toward the target. While using the Squeezer, subjects had no choice but to rely on this magnitude-only strategy to find the target angle, and it is notable that all subjects mastered it, achieving rise times and settling times that did not differ significantly from any other feedback types. Clearly, users are able to take advantage of the way tactile magnitude cues (intensity or frequency) vary over space when direction is difficult to discern or completely absent.

6.3 Trajectory Following

This final task required subjects to apply the tactile cue interpretation skills they had developed during angle targeting to follow a continually changing target. This task is significantly more challenging because the feedback changes even if the subject does not rotate his or her wrist. As with the previous task, several subjects occasionally felt lost at some point during a trial and became particularly flustered by a target that moved farther away even as they moved towards it. Since it punishes large deviations like these, our RMS error results indicate that Tapper Pulsing and Dragger Steady achieved the most

success in helping users track the trajectory. While Tapper Pulsing performed well on both other tasks, Dragger Steady yielded slow reaction times and long settling times, suggesting these simple metrics don't completely describe the potential of a certain form of tactile motion guidance.

Looking deeper into the results, we see that the Tapper, Dragger, and Vibration devices each have a drive algorithm that achieved significantly lower RMS error in trajectory following than one or more of the algorithms for the Squeezer or the Twister. This good quantitative performance is nicely reflected by the questionnaire data, wherein the Tapper, Dragger, and Vibration devices were perceived to be significantly more intuitive for the studied motion guidance tasks than the Twister and the Squeezer. The same three devices (Tapper, Dragger, and Vibration) were also rated to deliver direction cues that are significantly easier to discern, which may be an important key to perceived intuitiveness. In contrast, the ease with which cue strength can be discerned may not be as important to overall performance, as the Squeezer device and the Pulsing algorithms did significantly better in this category without good overall performance.

All ten subjects learned to interpret and react to the wide variety of tactile motion guidance cues tested in this study, even though they had no visual or auditory feedback. We were most surprised by the non-directional Squeezer's ability to guide users along trajectories with RMS errors that are comparable to most of the other devices. We largely attribute this finding to the fact that users could very precisely interpret the grading of the cue as the Squeezer loosened and tightened around the wrist, as evidenced by its high questionnaire ratings for ease of determining strength. Even though subjects had to initially guess which direction to move with the Squeezer, they could quickly recover from an error based on the grading of the feedback. Despite the relative success of this strategy, users did not rate the Squeezer very high for intuitiveness, indicating that they did not favor the initial guessing process. From these observations we conclude that humans are remarkably capable at interpreting a wide range of tactile cues in motion guidance tasks, even ones without direction information, but they prefer systems that clearly indicate both direction and magnitude.

7 CONCLUSION

This article presented a set of five wearable tactile actuators and two types of drive algorithms designed for guidance of wrist pronation/supination. Human subjects tested each actuator-algorithm combination in a set of directional response, angle-targeting, and trajectory-following tasks. Statistical analysis of the results showed significant differences between the ten forms of feedback for many of the tested performance metrics. This over-arching finding confirms our

hypothesis that both the selected actuator and the programmed algorithm play important roles in the success of tactile motion guidance systems.

The optimal combination of actuator and drive algorithm depends on the specific type of movement the user needs to make. Not surprisingly, subjects respond most quickly when they do not have to decide between two movement directions. When presented with directional tactile stimuli, they move most quickly when the cue's direction is conveyed through the location of the tactile stimulus (Tapper Steady, Tapper Pulsing, and Vibration Steady) or steady lateral skin stretch (Twister Steady). Feedback types that rely on a spatial tactile pattern (Dragger Steady, Dragger Pulsing, and Vibration Pulsing) elicit slower responses, as do Pulsing cues in general. The feedback types with the fastest directional responses also fare well in angle targeting, with Vibration Steady obtaining settling times that are significantly faster than the feedback types that use spatial patterns. Interestingly, subjects can use tactile magnitude cues (steady intensity and especially pulsing frequency) when direction is difficult or impossible to discern, as with the Squeezer. Finally, subjects follow arbitrary trajectories most accurately (median RMS error of less than ten degrees) with tactile feedback that is subjectively rated as "very easy" for discerning both movement direction and cue intensity (Tapper Pulsing and Dragger Steady), indicating the shared importance of these dimensions.

The best overall performance was achieved by Tapper Pulsing; this form of feedback certainly deserves additional study. Furthermore, many other actuator and algorithm options are possible beyond those tested here. We hope that other researchers working on tactile motion guidance will benefit from knowing the findings of this experiment, as it touches on several additional feedback characteristics that may be useful for a wide range of applications. In particular, our focus on solving the challenge of tactile guidance for wrist rotation led us to design simple one-degree-of-freedom tasks. To enable applications such as sports and dance training, future research should examine tactile motion guidance in multiple-degree-of-freedom tasks and more challenging ambient environments. We think such tactile cues may combine well with robotic tools for motor learning, and we are also excited about the potential use of low-cost tactile motion guidance systems to enable athletes and patients to practice movements at home.

ACKNOWLEDGMENTS

The authors thank Richard Stanley for helping to design the reported statistical analysis methods and guiding their implementation. We thank Karlin Bark for suggestions on the design of the tactile actuators and the experiment, and we thank Laurel Buxbaum and Steve Jax for insights on the potential of tactile

motion guidance in stroke rehabilitation. This material is based upon work supported by the National Science Foundation under grant 0915560. The final months of this work were also supported by the National Science Foundation Graduate Research Fellowship under award ID 2011124025.

REFERENCES

- [1] A. Bloomfield and N. I. Badler, "Virtual training via vibrotactile arrays," *Presence: Teleoperators and Virtual Environments*, vol. 17, no. 2, pp. 103–120, 2008.
- [2] J. Lieberman and C. Breazeal, "TIKL: Development of a wearable vibrotactile feedback suit for improved human motor learning," *IEEE Trans. Robotics*, vol. 23, no. 5, pp. 919–926, 2007.
- [3] B. Weber, S. Schätzle, T. Hulin, C. Preusche, and B. Deml, "Evaluation of a vibrotactile feedback device for spatial guidance," in *Proc. IEEE World Haptics Conf.*, 2011, pp. 349–354.
- [4] P. Kapur, M. Jensen, L. J. Buxbaum, S. A. Jax, and K. J. Kuchenbecker, "Spatially distributed tactile feedback for kinesthetic motion guidance," in *Proc. IEEE Haptics Symposium*, 2010, pp. 519–526.
- [5] K. Bark, P. Khanna, R. Irwin, P. Kapur, S. A. Jax, L. J. Buxbaum, and K. J. Kuchenbecker, "Lessons in using vibrotactile feedback to guide fast arm motions," in *Proc. IEEE World Haptics Conf.*, 2011, pp. 355–360.
- [6] K. L. Lurie, P. B. Shull, K. F. Nesbitt, and M. Cutkosky, "Informing haptic feedback design for gait retraining," in *Proc. IEEE World Haptics Conf.*, June 2011, pp. 19–24.
- [7] L. A. Jones and N. B. Sarter, "Tactile displays: Guidance for their design and application," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 50, no. 1, pp. 90–111, 2008.
- [8] M. A. Baumann, K. E. MacLean, T. W. Hazelton, and A. McKay, "Emulating human attention-getting practices with wearable haptics," in *Proc. IEEE Haptics Symposium*, 2010, pp. 149–156.
- [9] R. Tadakuma and R. D. Howe, "A whole-arm tactile display system," in *Proc. IEEE World Haptics Conf.*, 2009, pp. 446–451.
- [10] K. Bark, J. Wheeler, P. Shull, J. Savall, and M. Cutkosky, "Rotational skin stretch feedback: a wearable haptic display for motion," *IEEE Trans. Haptics*, vol. 3, no. 3, pp. 166–176, 2010.
- [11] B. Gleeson, S. Horschel, and W. Provancher, "Design of a fingertip-mounted tactile display with tangential skin displacement feedback," *IEEE Trans. Haptics*, vol. 3, no. 4, pp. 297–301, 2010.
- [12] A. A. Stanley and K. J. Kuchenbecker, "Design of body-grounded tactile actuators for playback of human physical contact," in *Proc. IEEE World Haptics Conf.*, 2011, pp. 563–568.
- [13] J. R. Winger, H. Burton, R. J. Sinclair, J. E. Brunstrom, and D. L. Damiano, "Joint-position sense and kinesthesia in cerebral palsy," *Archives of Physical Medicine and Rehabilitation*, vol. 90, no. 3, pp. 447–453, 2009.
- [14] R. Whelan, "Effective analysis of reaction time data," *The Psychological Record*, vol. 58, pp. 475–482, 2008.
- [15] U.-D. Reips and F. Funke, "Interval-level measurement with visual analogue scales in Internet-based research: VAS Generator," *Behavior Research Methods*, vol. 40, no. 3, pp. 699–704, 2008.



Katherine J. Kuchenbecker (S'04-M'06) is the Skirkanich Assistant Professor of Innovation in Mechanical Engineering and Applied Mechanics at the University of Pennsylvania. Her research centers on the design and control of haptic interfaces, and she directs the Penn Haptics Group, which is part of the GRASP Robotics Lab. Dr. Kuchenbecker serves on the program committee for the IEEE Haptics Symposium and other conferences in the field, and she has won several

awards for her research, including an NSF CAREER award in 2009, inclusion in the Popular Science Brilliant 10 in 2010, and a tie for Best Hands-On Demonstration at the IEEE Haptics Symposium in 2012. Prior to becoming a professor, she was a postdoctoral researcher at the Johns Hopkins University, and she earned her Ph.D. in Mechanical Engineering from Stanford University in 2006.



Andrew A. Stanley (S'10) received the BSE degree in Mechanical Engineering and Applied Mechanics from the University of Pennsylvania in 2011. His senior design team created a haptic CPR manikin that won the Gold award in the undergraduate division of the 2011 James F. Lincoln Arc Welding Foundation Engineering Student Design Competition. He worked in Dr. Kuchenbecker's Haptics Research Group until September 2011 and is now pursuing his graduate studies

in Mechanical Engineering at Stanford University, supported by a National Science Foundation Graduate Research Fellowship.