A glove for rubbing and touch gestures

Sam A. Miller, Andy Smith, and Robert St. Amant
North Carolina State University
Raleigh, NC 27695

July 12, 2011

Abstract

This paper describes the design and preliminary evaluation of a glove with individual sensors covering the tips of the thumb, third, and fourth fingers, plus grids of smaller sensors lining the palmar sides of the first and second fingers. The glove senses two types of gestures: rubbing the thumb against the first two fingers, and tapping the thumb on the tip of any of the four fingers, alone or simultaneously in different combinations. The glove occupies a novel point in the space of input devices, supporting low-resolution spatial input in one or two dimensions by rubbing, for discrete input by finger tapping, and for discrete input by single-stroke directional rubbing.

1 INTRODUCTION

This paper describes the design and preliminary evaluation of a glove instrumented with conductive threads and fabric. The glove contains large individual sensors covering the tips of the thumb, third, and fourth fingers; grids of smaller sensors line the palmar sides of the first and second fingers. The sensors of the first through fourth fingers are activated by contact with the thumb. Two types of gestures are supported: a touch contact between the thumb and the other fingers, and a rubbing gesture of the thumb on the first or second finger, or both together. Conceptually, the glove can be thought of as a low-resolution trackpad plus four buttons, all used by the thumb. Figure 1 shows the glove and the gestures it supports.

The glove shares interaction properties with some other glove-based input techniques. No object is held in the hand, which facilitates switching to other kinds of tasks that require use of the hand, allowing for the physical dexterity constraints that any glove imposes. Based on the taxonomic criteria of Dipietro et al. [4], the glove can be described as follows: it relies on 1 cloth-supported discrete sensor on each of three fingers and > 1 discrete sensors on two fingers, with 1-bit precision per sensor, and it is a tethered device (wireless is feasible but not currently implemented), communicating over a USB interface. The glove supports multiple-finger, surface-free input.

The glove integrates different interaction techniques in a novel way. First, input methods, including those associated with most types of gloves, are typically categorized as being continuous or discrete, with continuous input being used for spatial information in one, two, or three dimensions and discrete input for symbolic commands or characters. This distinction is not driven by physical or logical properties of interaction [2], but there are few common examples of linear input devices with an explicitly limited resolution used in computing today (e.g., such as a ratcheted physical slider). The glove provides spatial input, but at a much lower resolution than is typically associated with continuous input devices, and the sensors it relies on are simple discrete switches. Second, the glove exploits this low resolution input by allowing users to take advantage of proprioception, their implicit knowledge of locations on the fingers that can be touched by the thumb. Third, these touching actions are open-loop gestures that can be used for spatial targeting, which in most other input devices is supported by closed-loop gestures (e.g.,

a. The glove

b. Rubbing

c. Single click

d. Multiple-finger click

Figure 1: The glove.

targeting actions associated with pointer movements). These three interaction properties are not unique in isolation, but to our knowledge they have not been combined in a single input device.

The body of this paper describes work related to the design of the glove, then describes a prototype application in which the glove is used as an input device. Results from an experiment exploring the performance characteristics of the glove are given. The paper ends with a discussion of the current status of the project and future work.

2 RELATED WORK

Research on conductive and "smart" fabrics for HCI goes back over a decade (e.g. [13]); in recent years it has become relatively easy to develop and experiment with different uses of the technology. Beyond the basic medium of conductive threads and fabric, the glove draws on two main areas of work: glove-based interaction and the use of proprioception in input devices.

Dipietro et al. [4] extensively survey the literature of glove-based interaction, including glove character-

Figure 2: Robot views in simulation.

istics, application areas, and advantages/disadvantages. None of the 32 gloves surveyed has the combination of interaction properties of our glove; we believe that low-resolution spatial input by open-loop gestures is unique in glove-based interaction. The most similar device in the literature is the chording glove, which includes fingertip sensors and three discrete switches along the index finger [14]. The input device taxonomy of Card et al. [2] shows a scale for the resolution of input information via linear devices (e.g., the mouse), but almost all linear devices (excepting menus) provide "infinite" resolution in multiple dimensions (i.e., input that matches the high resolution of a display device). Ordinarily we would not see reduced resolution as an advantage in an input device. In the case of our glove, research in this direction is driven by two factors. First, low resolution facilitates a relatively natural interaction technique by exploiting proprioception. Second, conventional technology for high-resolution touch devices does not lend itself to this style of interaction, whereas low-resolution fabric-based technology does.

Proprioception is is an obvious factor in non-spatial input (e.g., in touch typing) and in targeted movement tasks, in both the directional component (e.g., movement of a joystick in one of the cardinal directions) and the distance component. Proprioception has been applied directly in the development of some input techniques, though to different ends than with our glove. Mine et al. [8] describe a framework for virtual reality systems in which a number of proprioception-based interaction techniques are introduced. De Boeck et al. [3], also working in virtual reality, describe techniques for analyzing the integration of proprioceptive cues into a multimodal interaction. These techniques apply at the arm- and body-scale, and they are not associated with tactile or haptic feedback.

A few projects described in the HCI literature are similar in spirit to our own, though they differ significantly in focus and scope. For example, Olwal et al. [10] demonstrate the promise of rubbing and tapping gestures for precise selection tasks on mobile devices. Li et al.'s soundTouch [5] supports vibrotactile interaction in the form of taps and rubs, but for output rather than input. In addition to performance findings, Li et al. report that subjects find the sensations from the device to be similar to touches from other humans. Ni and Baudisch [9] describe the evolution of mobile devices toward disappearance via miniaturization. They evaluate the use of touch and motion scanners for character entry and marking tasks. We expect that some categories of user errors they observed, dealing with device alignment and detection range, would be reduced in our glove—though as we will see, the glove has other limitations.

There is also a long history of gloves for automated recognition of sign language (Ong and Ranganath [11] give a relatively recent partial survey).

3 THE GLOVE

The glove consists of a set of digital and analog data channels in a combination of conductive thread and wire. The glove can be thought of as a set of switches, each of which completes a circuit through a a variable resistor, which is an individual conductive thread. A switch is closed when the thumb makes contact with any thread. Current is supplied by a National Instruments USB-6009 device.

The grid, divided into two even parts over the first and second fingers, provides 80 sensed locations in all: four threads running the length of each finger (the y threads), ten across the width (the x threads). The threads are sewn in a crosshatch pattern on and below the surface of the glove fabric, such that the threads do not touch each other, in particular at crossing points in the x and y dimensions. When the thumb touches the grid, it contacts at least one x and one y thread, closing those circuits. Each such event generates a bounding box (or a single point) on the grid, from which the center and the size of the box can be extracted, as shown in Figure 3.

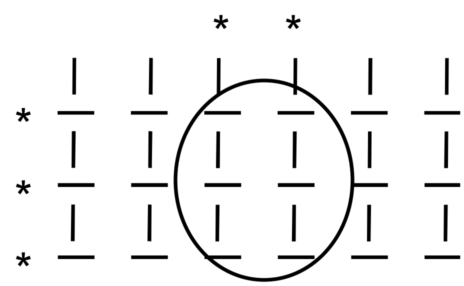

Figure 3: Circuits activated by thumb on glove.

If the bounding box covers an odd number of grid lines in a given dimension, the center of the box corresponds to the center grid line in that dimension. If the bounding box covers an even number of grid lines, however, a point between the two center lines is used. This gives the grid an effective resolution of almost quadruple the number of its sensed locations (with the qualifier being due to locations on the boundaries). This provides the intuition behind the glove's input processing. The actual implementation treats the grid as pixels of an image. Through multiple steps involving super resolution, look-up tables, and morphological reductions, the energy center of the touch "image" is extracted. These steps significantly reduce the noise inherent to the thumb/finger contact point, leverage the fact that the thumb contact area is usually quite large, and guarantee a single grid coordinate as the output value. This last point is notable as the grid intrinsically facilitates multi-touch interactions similar to those normally found in modern touch screens, but because the thumb is the only contact point with the grid, a single output is desirable.

At an abstract level, the glove can be compared with other input devices in the form of a three-state model [1], as shown in Figure 4. The size of the bounding box can be treated as further input, but we have not explored the use of this extra information in our work to date.

4 AN APPLICATION

The initial goal for the development of the glove was to support flexible one-handed interaction with a synthetic display, as might be provided in an augmented reality system, while the user is walking and looking around through a head-mounted display (HMD). A simulation of such an application was developed, based on a video game engine. The simulation contains two semi-autonomous robots, a crawler and a flyer, that follow pre-programmed paths through a rain forest. Each robot pipes its video feed and other status information to the user. The user can take control of either machine at any time to direct its actions using buttons, menus, and navigational controls.

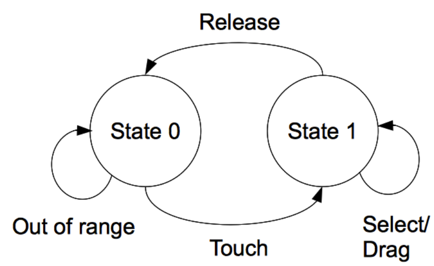

Figure 4: Three-state model for the glove.

The simulation is operated by the glove via different types of gestures:

- Some global commands are mapped to taps on fingertips: a single tap on one or more fingers, or repeated taps on one or more fingers. We call this a *tap*.
- Some local or context-sensitive commands are mapped to directional rubbing gestures, single strokes on the first or second finger, or across both. These can act as shortcuts to execute a command or change state in the interface (a *symbol rub* gesture).
- Advancing down linear menus, a spatial command in one dimension, is mapped to a rubbing gesture on the first finger (a *scroll* gesture). This is distinguished from symbol rub gestures in that it involve targeting by distance as well as direction. (Scanning with the camera would also be possible with this gesture, though it is not implemented.)
- Continuous, dynamic directions in two dimensions, analogous to the output of a joystick, are supported by rubbing over the first two fingers (a *2D rub* gesture). An analogous use of this gesture in the application is for choosing from a pie menu.

These gestures can be seen in a sample scenario in the simulation, in which the user responds to a call from the crawler that is stuck and needs assistance. The user executes the following steps, with the category of the relevant gesture given in italics.

1. Bring up a standard view on the HMD, including a diagnostic window. *Tap.*
2. Switch to the crawler's control interface. *Tap.*
3. Shift focus to the "Behaviors" menu. *2D rub.*
4. Enter the "Behaviors" menu. *Tap.*
5. Move through the items to highlight the "Manual" behavior. *Scroll.*
6. Select the behavior. *Tap.*
7. Switch focus to the "Drive" control. *2D rub.*
8. Start manual drive. *Tap.*
9. Drive the crawler until it is free. *2D rub.*

Formative usability tests were conducted with six users. Each session began with a demonstration to show the different interaction options the glove supports. Within five minutes all users were reasonably competent in the interface, the only consistent challenge being the difficulty of remembering the mapping of different gestures to commands. Qualitative feedback was positive: Users found some kinds of gestures intuitive and enjoyable, in particular control of the movement of the robots by rubbing gestures.

A few relatively obvious lessons came to light during the formative evaluation. One is the potential difficulty for novice users to learn the mapping of gestures to commands in cases where there is no intuitively obvious choice, such as changing modes by tapping all fingers at once. (This was gesture was used to avoid accidental mode changes.) Another is the value of designing the layout of the interface such that, in specific modes, interface components can be selected by a short directional rub. This raises the issue, however, of whether strongly modal interfaces would be appropriate in other task domains.

The glove generally met our expectations in the application. It appears to show promise in offering an effective way to control a specific simulation, with some commands and continuous actions naturally mapping to the gestures supported by the device.

5 EVALUATION

We have carried out an informal evaluation to test the performance characteristics of the glove in tasks for which we expect it to be well suited. Our interest is in the following:

- *Target selection.* Consider a scenario comparable to the one described in the previous section, in which the user would like to select a target region of a pre-specified size, perhaps indicated on a head-mounted display, using the glove. How accurately can the user perform such an action, either without visual feedback for the action or with adjustment based on feedback?
- *Fingertip tapping gestures.* The glove supports discrete, non-spatial actions in the form of tapping the thumb to different fingers. How quickly and accurately can users carry out such actions?
- *Directional rubbing gestures for symbolic input.* Directional rubbing gestures on the glove can interpreted as discrete symbols. How quickly and accurately can users carry out sequences of such actions?

The goal of our evaluation was to develop a set of quantitative performance benchmarks for the use of the glove under laboratory conditions, mainly in the form of descriptive statistics. Tasks were designed to answer the questions above, and a testing environment was created for each, with parameters selected based on a pilot study with the one of the authors acting as a participant.

We discovered that we had underestimated the noisiness of the prototype as an input device, due to its physical construction, its sensor limitations, and our input processing techniques. For example, the glove must be positioned appropriately on the fingers to function properly, and it does not fit all hands. In use, the glove sometimes registers spurious contact/release events, and one or two threads on the glove occasionally miss input, due to bending strain on the connections. These issues play a role in limiting performance in all the tasks described below.

We carried out an experiment involving four tasks, with 11 participants, all unpaid volunteers. Participants ranged in age from 21 to 33. Two were female, nine male. All were right-handed, with vision corrected to 20/20, and use computers daily. The total duration for a session was about 30 minutes per participant. Each participant first put on a non-latex insulating glove and then the haptic glove. Tasks were carried out in the same order for all participants. Between tasks, participants were allowed to rest for up to two minutes.

Tasks 1a and 1b: Indicating spatial location (Part 1). Task 1 was a target selection task in one (Task 1a) or two (Task 1b) dimensions. Task 1a was performed before Task 1b. Participants were shown a box containing a colored target region, as in Figure 5.

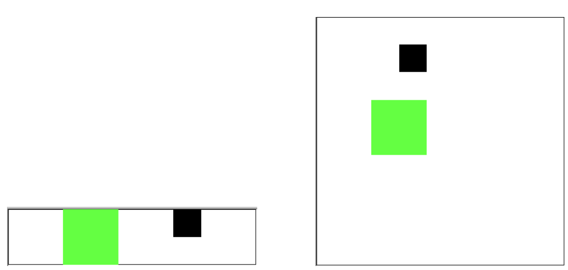

Figure 5: Target selection in one and two dimensions.

The spatial extent of the box is mapped in absolute terms to the layout of sensors on the first two fingers of the glove. In Task 1a, vertical movement is ignored. The box is divided into an array of $n \times 1$ elements for Task 1a or $n \times n$ elements for Task 1b (the resolution), and the target occupies s adjacent elements for Task 1a, $s \times s$ elements for Task 1b (the relative target size). These divisions are not visible in the display. The specific values used for n and s are given in Table 1, based on our judgments about the difficulty of the task from the pilot study.

Note that $n = 30$ is an interesting challenge for the glove, which works at a resolution of at most 10 discrete sensor locations in one dimension. Unlike mouse interactions, the mapping from the input space to the output space is one-to-one and onto. That is, we are asking participants to control a space with a higher resolution than that of the input device.

n	s	*Ratio*
9	2	0.22
9	4	0.44
30	2	0.06
30	4	0.13
30	8	0.26

Table 1: Resolution (n), target size (s), and target size ratio.

For Task 1a and then for Task 1b, the participant carried out 24 trials in a training phase. Participants then carried out 10 trials for each of the conditions in Table 1, with target locations generated at random by the system. A trial began with the participant's thumb moving freely, not in contact with the sensors on the fingers. A target appeared in a random location, and the participant touched the thumb to the fingers, causing a cursor to appear. The participant's goal was for the cursor to overlap the target. If this did not happen on the initial touch, the participant slid the cursor by rubbing with the thumb until the target was reached. The participant released contact with the thumb to indicate completion of the trial, and the next trial began immediately. For each trial, the duration was recorded between the release event at the end of one trial and the release event of the next trial; the time for a new target to appear was considered negligible. A trial was considered successful if the release event occurred with the cursor over the target.

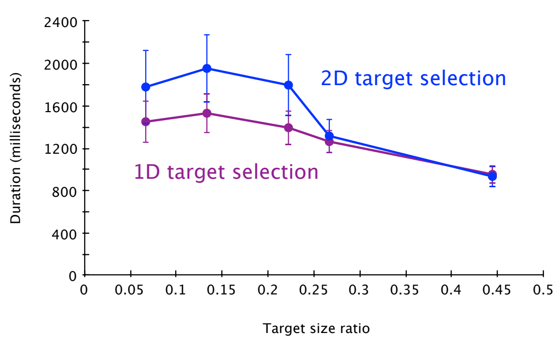

Figure 6: Task 1a/b selection duration, per target size ratio.

Some data cleaning was performed to remove outliers with durations longer than 10 seconds (less than 1% of the data for Task 1a, less than 2% of the data for Task 1b). Qualitatively, the results were largely as expected. Figure 6 shows the change in duration of target selection as the target size ratio increases. Figure 7 shows the frequency of successful completion of a trial.

Over all the conditions in Task 1a, the mean duration for target selection was 1326 milliseconds. The fastest condition was a resolution of 9 and a target size of 4 (for a target size ratio of 0.444), with a mean duration of 953 milliseconds. The most successful condition was also 9/4, with a success frequency of 0.82. The lowest success frequency by an individual participant was 0.071, in condition 30/2; no

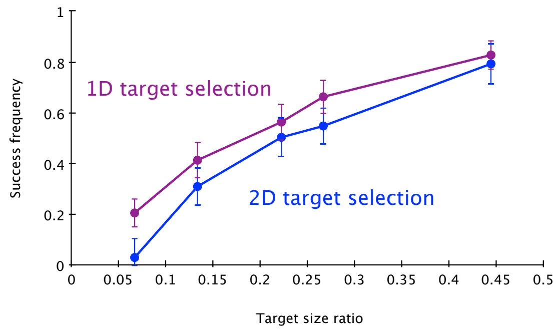

Figure 7: Task 1a/b success frequency, per target size ratio.

participant reached a frequency higher than 0.50 for that condition. The highest success frequency by an individual participant was 1.0, reached by seven participants in various conditions (9/4, 30/8, and 30/4).

Task 1b values are comparable: the overall mean was 1562 milliseconds, the best condition (9/4) at 934 milliseconds, with the success frequencies being 0.56 and 0.79, respectively. Eight participants were unable to score higher than 0 in the 30/2 condition; four participants reached a score of 1.0 in the 9/4 condition.

For context, Parhi et al. [12] give duration results for selection of targets on small touchscreens with the thumb, in one-handed use. For targets of average size on small touch screens, selection time is between 1100 and 1200 milliseconds; errors appear in almost 30% of selections. The tasks in their experiments are significantly different from Tasks 1a and 1b, but the comparable timing and error rates (e.g., for the midpoint in Figures 6 and 7) suggest that glove interaction is not unreasonably slow. The low success rate of 0.82 for even the best target size ratio, however, is more problematic. As mentioned above, this is attributable to two factors: a high level of noise in the sensing for the glove (i.e., spurious release events detected), and the limitations the glove imposes on fine targeting actions.

Task 2: Indicating spatial location (Part 2). Task 2 was designed to test whether Fitts' Law applies to the selection of targets only using rubbing gestures, patterned after MacKenzie and Buxton's study of two-dimensional pointing tasks [6]. The structure of the task was similar to that of Tasks 1a and 1b, with the display functioning as a slider widget. Participants first acquired the indicator of the slider and then, by a rubbing action, dragged it to a specified target location, again releasing the thumb at the end of each trial. Unfortunately, this proved unexpectedly difficult for two reasons. Some of the target locations were so close to the boundaries of the device that they could not be reached with a continuous rubbing action, and the initial acquisition of the slider indicator proved more difficult than we had anticipated. For more than half of the participants, the task quickly became so frustrating that they gave up.

We do not report results for this task; however, its failure under experimental conditions suggests the limitations of the device. The current design, in which input and output are decoupled with respect to resolution and a confirmation action simply involves release of the thumb, is inadequate for interaction with conventional graphical user interfaces. Obvious solutions include confirmation actions that are separate from movement, as in the simulation application of the previous section, but new interaction techniques may be needed. Qualitatively satisfactory results in the robot application, where pie menus and taps were used extensively, indicate those modes of interaction better leverage of the device's capability, as contrasted with the poor performance for other types of pointing operations.

Tasks 3a and 3b: Providing discrete input with finger taps. Task 3a involved participants "playing the scales" on their fingers. After a practice phase of 24 trials, participants executed 50 trials, tapping the first through fourth finger in order, then in reverse order. Only duration was measured for these trials, from one release event to the next. The mean duration per tap was 305 milliseconds.

Task 3b required participants to watch the computer display and follow instructions for which finger to tap. Instructions were presented visually, in the form of four colored squares in a vertical column, each corresponding to a different finger. A trial consisted of one square turning a different color and the participant tapping the appropriate finger. The next trial began immediately thereafter. An error was counted if a participant touched a finger other than the one shown on the display; the participant would repeat the trial with the correct gesture. In a practice phase, participants executed 24 trials using the interface, and then carried out 50 trials for the task, in which the finger to tap was determined randomly. Duration and success (correct taps) were recorded.

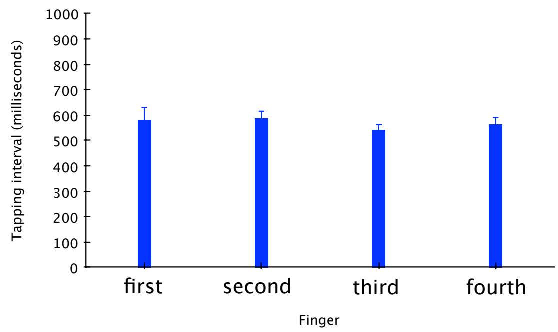

Figure 8: Task 3b tap interval, per finger.

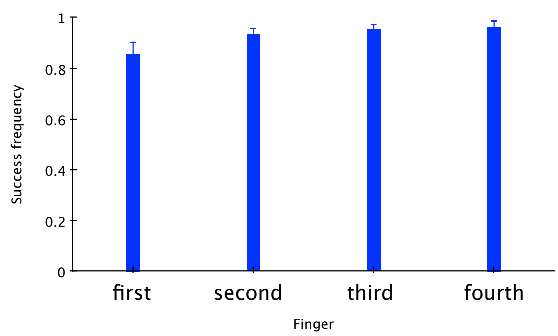

Figure 9: Task 3b success frequency, per finger.

Two outlier trials were removed before the analysis was carried out. Figures 8 and 9 show the duration and success frequency per finger. The mean over all fingers was 575 milliseconds, with a success frequency of 0.93. The higher duration and lower accuracy of tapping the first finger (the index finger) goes against expectations; we believe it should be attributed to the glove hardware and input processing. These durations are in the expected range for the glove, comparable to performance on the chording glove [14], though error rates appear to be higher.

Task 4: Providing discrete input with directional rubbing. Task 4 is a gesture entry task, analogous to text entry tasks in which participants enter gestures using an input device, following a sequence provided on a computer display [7]. All gestures in Task 4 are symbol rubs, as described above, single strokes in a diagonal direction. A sampling is shown in Figure 10. The direction of the line segment indicates the direction of the gesture, and the red dot indicates the starting point. The upper box indicates a symbol rub on the first finger, the lower box on the second finger; a tall vertical box indicates a symbol rub across across both fingers. (Note that in this task, participants are not asked to remember the mapping of gestures to symbols, which means that timing and errors are due to the text copying process rather than memory.)

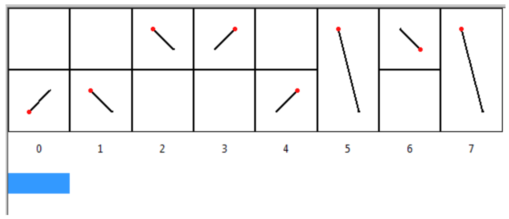

Figure 10: Symbol rub gestures.

A set of trials began when the participant tapped the thumb to the fourth finger; participants carried out the sequence of gestures displayed. A blue cursor advances through the displayed gestures as they are carried out. On reaching the end of a sequence, participants tapped the fourth finger to advance to the next sequence.

Participants carried out 24 trials to practice the gestures. During the practice phase, a simple gesture recognition system provided visual guidance as to the accuracy of the executed gesture. Gesture recognition was turned off for the trials that followed, because we believed that this would provide for more realistic estimates of the speed of gesture entry in practice, and because of the unreliability of input processing. After the practice phase, participants carried out 50 gestures, following a randomized sequence, broken up into sets of 8, as shown in Figure 10.

The mean duration of a gesture in Task 4 was 364 milliseconds, and the success frequency (i.e., recognition rate) was just 0.40. The relationship between duration and recognition rate is show in Figure 11. The lowest success frequencies per participant (0.02, 0.06, and 0.08) are associated with gesture durations of 200 milliseconds or less. The highest success frequency, 0.94 for one participant, was associated with a mean duration of 614 milliseconds per gesture. The mean 40% recognition rate is clearly far too low for practical use of the glove in this task. Based on our observations of the tests, it seems possible the low recognition rate may be partially due to the horizontal data lines that intermittently fail; noise in the form of spurious release events also played a role. A less noisy glove my improve results.

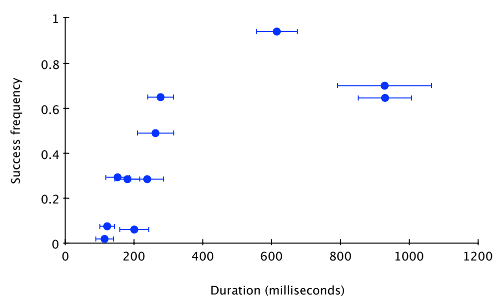

Figure 11: Task 4 duration and success frequency, per participant.

6 DISCUSSION

Contrary to our expectations, performance for target selection (Tasks 1a and 1b) and symbol input (Task 4) was surprisingly low. Worse, a mouse-style pointing task had such low performance that the task could not be completed. We partly attribute this to the glove's noisy capture of the user input.

Despite the disappointing overall results of our experiment, we find a few bright spots: a few participants, in a few tasks, achieved reasonable performance. The authors of this paper, with more experience using the glove, significantly outperform the best of the participants in the experiment; as in other glove-based work (e.g. [14]), training plays a factor in performance.

The simulated robot application, in contrast, worked much better. Failed inputs had low cost—actions could be repeated—and users found navigating to icons using a pie-menu interaction, menu selections with 1D scrolling, driving and flying the robots with joystick interactions, and shifting interface modes or confirming selection with taps, to all be highly effective, easy to learn, and not frustrating. This suggests that gloves that use a "touchpad on the finger" may have relevance for specific application domains and interaction styles, in particular those relevant to interaction when the user's hand might be out of view. We believe that with better hardware, the glove should offer an interesting new avenue for interaction in wearable computing.

Our plans for future work include the following:

- Exploring alternative sensing technology for the glove. One plausible candidate is circuitry printed on the fabric of the glove; another is fabric-based rheostats.
- Human performance experiments. With a more reliable glove, it should be possible to estimate the number of distinct locations users can distinguish when tapping or rubbing along the length of their fingers. Task 1 offers no more than a rough approximation of the ability to select different targets.
- Expansion of the gesture set for symbol rubs. We are interested in the possibility of text entry by glove gestures. It is straightforward to carry out distinct, single-stroke gestures with the thumb on the surface of the first two fingers without wearing a glove. We hope to bring such easy, natural movements to glove-based interaction.

Our last goal is exploring new possible applications for the glove. One of the most appealing aspects of the style of interaction is its expressiveness, in the same sense that gesture-based systems can support expressive actions. We have begun to build an application for generating music with the glove, based on instruments such as harmonicas that do not necessarily depend on precision for their effectiveness. Even in its present state, the glove offers some promise.

References

[1] W. Buxton. A three-state model of graphical input. In *Human-computer interaction—INTERACT '90*, volume 90, pages 449–456. Elsevier, 1990.

[2] S.K. Card, J.D. Mackinlay, and G.G. Robertson. The design space of input devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Empowering people*, pages 117–124. ACM, 1990.

[3] Joan De Boeck, Erwin Cuppens, Tom De Weyer, Chris Raymaekers, and Karin Coninx. Multisensory interaction metaphors with haptics and proprioception in virtual environments. In *Proceedings of the third Nordic conference on Human-computer interaction*, NordiCHI '04, pages 189–197, New York, NY, USA, 2004. ACM.

[4] L. Dipietro, A.M. Sabatini, and P. Dario. A survey of glove-based systems and their applications. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 38(4):461–482, 2008.

[5] Kevin A. Li, Patrick Baudisch, William G. Griswold, and James D. Hollan. Tapping and rubbing: exploring new dimensions of tactile feedback with voice coil

motors. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, UIST '08, pages 181–190, New York, NY, USA, 2008. ACM.

[6] I.S. MacKenzie and W. Buxton. Extending Fitts' law to two-dimensional tasks. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 219–226. ACM, 1992.

[7] I.S. MacKenzie and R.W. Soukoreff. Text entry for mobile computing: Models and methods, theory and practice. *Human-computer interaction*, 17(2):147–198, 2002.

[8] M.R. Mine, F.P. Brooks Jr, and C.H. Sequin. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 19–26. ACM Press/Addison-Wesley Publishing Co., 1997.

[9] T. Ni and P. Baudisch. Disappearing mobile devices. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 101–110. ACM, 2009.

[10] A. Olwal, S. Feiner, and S. Heyman. Rubbing and tapping for precise and rapid selection on touchscreen displays. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 295–304. ACM, 2008.

[11] S.C.W. Ong and S. Ranganath. Automatic sign language analysis: A survey and the future beyond lexical meaning. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(6):873–891, 2005.

[12] P. Parhi, A.K. Karlson, and B.B. Bederson. Target size study for one-handed thumb use on small touchscreen devices. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, pages 203–210. ACM, 2006.

[13] E.R. Post and M. Orth. Smart fabric, or "wearable clothing". In *First IEEE International Symposium on Wearable Computers*, pages 167–168. IEEE, 1997.

[14] R. Rosenberg and M. Slater. The chording glove: a glove-based text input device. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 29(2):186–191, 1999.