

Grabbing at an Angle: Menu Selection for Fabric Interfaces

Nur Al-huda Hamdan¹, Jeffrey R. Blum², Florian Heller¹

Ravi Kanth Kosuru¹, Jan Borchers¹

¹RWTH Aachen University, 52056 Aachen, Germany

²McGill University, Shared Reality Lab / Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT), Montréal, Québec, Canada

ABSTRACT

This paper investigates the pinch angle as a menu selection technique for two-dimensional foldable textile controllers. Based on the principles of marking menus, the selection of a menu item is performed by grabbing a fold at a specific angle, while changing value is performed by rolling the fold between the fingers. In a first experiment we determined an upper bound for the number of different angles users can reliably grab into a piece of fabric on their forearm. Our results show that users can, without looking at it, reliably grab fabric on their forearm with an average accuracy between 30° and 45°, which would provide up to six different menu options selectable with the initial pinch. In a second experiment, we show that our textile sensor, Grabrics, can detect fold angles at 45° spacing with up to 85% accuracy. Our studies also found that user performance and workload are independent of the fabric types that were tested.

CCS Concepts

•Human-centered computing → Interaction devices; Interaction techniques;

Keywords

Textile interfaces; smart fabric; wearable computing; marking menus; input device

1. INTRODUCTION

Fabric and clothes are a ubiquitous part of our everyday environment, but despite this, they are rarely used as an input surface for interactive systems. Commercial products mostly transfer known interaction concepts like buttons and sliders to the textile domain, and do not take advantage of the natural affordances of cloth. However, textile materials allow for much richer interaction including folding, stretching, draping, and crumpling, which can be used to increase the input interaction bandwidth for wearable devices.

For example, if we look at smart-watches, where input and output compete for the limited display real estate, augmenting nearby fabric to function as an input controller helps mitigate this dilemma.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. ISWC '16, September 12–16 2016, Heidelberg, Germany Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4460-9/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2971763.2971786>

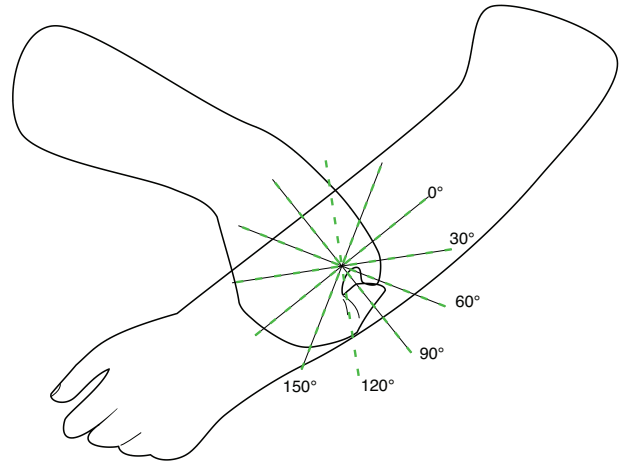


Figure 1: In our system, the user grabs and rolls a fold in the sleeve on her forearm. We investigated how well users could reproduce folds at different target angles as a means of determining which value to control with the fold.

In this paper we explore the possibilities of orientation-dependent fold-based interaction (Figure 1). The initial orientation, or angle, of the fold is used to select between the different values that can be changed, and similar to Pinstripe [12], the relative displacement when rolling the fold changes values. For example, pinching and rolling a fold parallel to the arm could change the volume, while pinching and rolling a fold perpendicular to the arm could skip items in a playlist. Note that this means that the orientation of the initial grab is absolute in reference to the user's body, while rolling is relative. Thus, the proposed input technique leverages the ability to interact wherever convenient on the body in the context of current activity, with the body itself serving as a reference and therefore taking advantage of proprioception.

This input technique was inspired by marking menus [14], a variant of pie menus, that in one mode allows the user to select a menu item by making a straight mark in the direction of the desired item without showing the menu itself. Empirical investigations [30, 2, 18] show that the concept of marking menus supports mobile and eyes-free interaction by using touch and motion as input. A general improvement to marking menus are Control Menus [23], which use the continued motion of an input device after an item has been selected to change a continuous value associated with that item. In this paper we apply these principles to textiles and use the fabric pinch angle and relative rolling of the fold as input. We implement these techniques using the Grabrics smart textile [6].

Further, pinching is an explicit gesture that can minimise accidental activation, e.g., when brushing against other surfaces, and enables easy and quick interaction. When triggered, using the pinch angle to immediately select a menu item takes away the need for individual physical controllers to access each item. This way, users would also be able to interact with fabric eyes-free, depending on their own body as a reference.

Previous investigations showed that users prefer, from both physical and social perspectives, to interact with fabric on the upper arm and forearm [12, 24]. Interacting with fabric on the arm is convenient for a wide variety of activities and postures, thus, we selected the forearm as the most practical site for our interface.

This paper makes three contributions. First, we report results showing that users with limited training can reliably grab fabric on their arm at angles separated by between 30° and 45° , which would provide up to 6 different menu options selectable with the initial pinch. Since this was carried out on swatches of fabric without any electronics built-in, we believe that this represents the limit of human performance in eyes-free pinching of fabric. Second, we show that this accuracy is possible on five different fabrics with varying characteristics. Third, we describe a user study to determine how people grab and fold a smart fabric such as Grabrics (described below), and discuss the limitations of the existing implementation in accurately distinguishing between different angles. Accordingly, we propose an alternative mechanism to earlier work [6] to improve the accuracy of fold angle detection in future implementations.

2. RELATED WORK

The question where to place an input device on the human body has come up repeatedly in the literature. Thomas et al. [27] attached a regular PC touchpad at different positions on the body and measured input performance in varying postures. The preferred location for the touchpad was the forearm, which is easily reachable under many circumstances, a finding confirmed by Karrer et al. [12]. Interacting with the forearm is also perceived to be socially acceptable in public contexts [24], a factor that is important for the general acceptance of such wearable controllers.

Fabric interfaces, especially in commercial products, often simply transfer known concepts to the textile domain, for example buttons [11] or touchpads [8]. Leveraging the textile nature of the materials allows to embed the sensors more seamlessly into our environment [21].

Grabbing and folding is part of our everyday interaction with fabrics. This motivated Pinstripe [12], which senses the size and relative movement of a fold that the user grabs into a piece of cloth, and gives her granular control over continuous linear values such as the volume of an MP3 player. However, Pinstripe is a one-axis sensor, which limits its input bandwidth and restricts the user to performing the fold parallel to the conductive stripes integrated into the garment. The Grabrics textile controller [6] can sense rolling gestures on the fabric in arbitrary axes.

But folding as a means of interaction can be much richer if we consider different types and directions of folds. For example, bending the edges and corners of a flexible display can be used to navigate in an e-book [3] and [10]. Khalilbeigi et al. [13] investigates a gesture alphabet for rigid, but foldable, displays whereas Lee et al. [15] defined a gesture alphabet of possible fold, bend, and distort gestures for paper, plastic, and stretchable fabric. How users naturally interact with fabric (e.g., pinch, stretch, squeeze, drape, etc.) has been recently motivated as an interaction metaphor for deformable user interfaces [28, 16, 22].

Detecting bends and folds around a single axis, such as the knee, has been shown to be possible with stitched sensors [5]. How-

ever, simultaneously sensing many different distortions applied to a piece of fabric is difficult, because stretching multiple axes requires many or very complicated sensors, which results in an overloaded fabric that is difficult to manipulate or wear in practice.

Capacitive touch detectors are an unobtrusive alternative to fabric interfaces for augmenting the body with input controls, as described in [17, 29] or with a sensing belt [4]. However, capacitive sensing has the drawback that you cannot search haptically for the touch surface because a touch event is triggered once you first make contact.

Researchers attempted to overcome the limited input space of mobile and wearable systems by designing gestural menus that are based on the concepts of marking menus [14]. For example, ear-Pod [30] allows users to access up to 8 items of a spatial audio menu eyes-free by sliding the thumb on a hand-held circular touchpad. Similarly, in [2] users could select from four items using head gestures, or select items using a pointing device [18].

The usage of marking menus for eyes-free interaction in a mobile context has been investigated using the motion sensors of a smartphone to detect tilt angles [20], as well as in combination with touch input [1]. This means we can rely on users memorizing the spatial layout of the values they can control using our sensor.

Our approach makes use of fold-based interaction and the pinch orientation and triggers only after a distinct activation by the user, i.e., grabbing a fold into the textile. Further, since our technique uses just the angle of the pinch to select an option, a user can choose between multiple items by pinching this single control at particular angles, rather than moving between multiple separate sensors (e.g., discrete buttons in different body locations) to select or control more than one value.

In the remainder of this paper, we first determine the number of reliably distinguishable angles a user can grab into a textile. Second, we demonstrate the feasibility of this approach based on a smart fabric prototype, Grabrics, and provide results of a user test and implementation of detecting the fold angle of the fabric.

3. EXPERIMENT 1

Human performance in fabric folding angles

The first experiment focused on establishing how many angles a user can reliably fold in fabric covering their forearm without looking at it, as a means of establishing an upper limit to the number of menu items that could be differentiated in a wearable marking menu-type interface.

3.1 Setup

We recruited eight participants (6 male, 2 female, age range 25-32, one left handed) from the university community. After completing a consent form and an initial survey of their demographic and physical data, including measurements of their hand and arm, each participant sat in a comfortable position with their arm resting on a table. A long, thin plastic cable tie (one participant), or a stiffer piece of thin spring steel (all other participants), was attached to their thumb, such that when they performed the grab, it was easy to see the direction of the pinch in the pictures during analysis (Figure 2). After being familiarized with the experimental procedure, the participant was shown an image of the ANGLE at which to grab the fabric on her forearm (Figure 2). Participants were instructed to not look at their arm while performing the grab, and the screen showing the target angle was located in a different direction from the arm wearing the fabric, forcing them to look away. In a training phase before the experiment began, using whatever clothing they happened to be wearing, they were given verbal feedback by the




Figure 2: The users were presented with a visual representation of the angle along which they were supposed to grab the fold into the fabric. The result was captured by a camera.

experimenter as to whether they had grabbed correctly given the presented image. After grabbing, they were asked to freeze for a brief period so a photograph could be taken using a GoPro camera mounted on a tripod directly above their arm. Since the eventual goal of the interface is to use the initial pinch axis to select an articular option, and then roll the fingers to change the value of that option, the participants were then asked to slide their fingers back and forth to simulate this, in order to make sure that their method of grabbing the fabric would support such a scenario.

Once they had demonstrated to the experimenter that they understood the task, the experiment began. A 25×25 cm section of one of five different fabrics was attached to their non-dominant forearm using a Velcro band just below the elbow, with the fabric loosely extending to approximately the participant’s wrist area (Figure 2). Each fabric was marked with a 0.5 cm grid to facilitate calculating the angle of the initial grab. Although not formally characterized, the five fabrics exhibited stiffness ranging from approximately denim-weight, to a more silky texture and drape (Table 1). All fabrics were pre-washed to remove textile finishes, such as Formaldehyde, which acts as a preserving coat on new fabric to reduce wrinkles and shrinkage.

Fabric	Specification
Jersey	96% cotton, 4% spandex
Knitted Jersey	93% cotton, 7% spandex
Stretch-Satin	97% polyester, 3% spandex
Poplin	35% cotton, 65% polyester
Jeans	100% cotton

Table 1: Fabrics tested in Experiment 1.

The experiment used a within subjects design. Participants interacted with each of the five fabrics one at a time, with the order counterbalanced. For each fabric, participants were asked to mimic each of 6 angles (From 0° – 150° in 30° increments), counterbalanced using a Latin square design within each block, since the experimenter observed in an earlier pilot study that having them in a consistently increasing order resulted in a learning effect. This block was repeated three times, then the next fabric was placed on the participant’s arm. After each block of trials on each fabric, the participants were asked to answer three 5-Likert scale questions on the easiness of pinching, comfort of pinching, and easiness of rolling. With five fabrics, six angles, and three repetitions, each participant completed 90 total trials.

3.2 Results

In addition to the experiment described above, a pilot study with a similar setup was conducted beforehand with eight participants (6 male, 2 female, age range 22-28, one left-handed, none of whom

participated in the first experiment). This pilot established an appropriate number of repetitions without user fatigue and determined which fabrics would be used in the study. In addition to the steps described for the main experiment, participants also completed a NASA TLX [7] questionnaire after each fabric. We mention this pilot in order to report that the fabric type had no effect on the NASA TLX scores, all of which were less than 50%, indicating that no significant issues with mental or physical workload were evident from using the selected fabrics. Last, the first pilot participant also used a sixth leather-like fabric, which was removed from the main experiment since he could not effectively grab it to form a pinch. Thus, we conclude that some fabrics are clearly not suitable for pinch interfaces due to their physical properties. The NASA TLX portion was not repeated in the main experiment in order to reduce the total time.

Once the data was gathered from the main experiment, the fold angles were measured using the ruler tool of a photo editing application. The measurements could be reproduced independently by two different experimenters to an error of around 3°. From the entire set, 15 images were too blurry to be usable, and were excluded from analysis.

A repeated measures ANOVA on log-transformed absolute errors with user as random factor found a significant main effect of ANGLE on the error ($F_{(5,694)} = 10.981, p < .0001$). A closer look at the results with a post-hoc Tukey HSD test shows that the error for the fold along the arm (0°) is significantly ($p < .001$) lower than for all intermediate angles (30°, 60°, 120°, 150°), but not compared to a fold perpendicular to the arm (90°). Although the average error (Table 2) indicates that a 30° spacing as in our experiment could work, the distribution in both directions results in overlap that would make a clear separation of the users’ intentions difficult (Figure 3). We therefore recommend using a separation of at least 45°. The type of FABRIC did not have a significant effect on the measured error. Last, we note that with longer training time, performance may improve. This may be even more likely if the user interface provides feedback on the accuracy of the folding, thus training the user to modify their folds to make them less ambiguous to a classifier.

		Target Angle					
		0°	30°	60°	90°	120°	150°
Error	M	5.0	8.6	10.6	6.7	14.2	10.9
	SD	4.4	6.1	9.8	4.7	9.23	8.8

Table 2: Average measured errors by target angle.

We used ordinal logistic regression to evaluate how the angle and fabric affected participant response to our 5-Likert scale questions. We found only a significant effect for ANGLE on the easiness of pinching ($\chi^2_{(5, N=240)} = 30.0, p < .0001$) and general comfort ($\chi^2_{(5, N=240)} = 50.4, p < .0001$). We found a significant effect for FABRIC on the easiness of rolling the fabric after pinching ($\chi^2_{(5, N=240)} = 36.7, p < .0001$). From experience with various prototypes [26], the texture of the inner surface of the sensor is of similar importance to the top layer, since the friction of the fingers on the fabric must be greater than the friction of the inside of the fabric rubbing on itself in order to roll the fold.




Figure 3: Measured vs. specified angles. We see that the spacing of 30° is at the limit of what users can reliably distinguish.

4. EXPERIMENT 2

Grabrics fold angle test

Although the above experiment uses only normal fabrics, without any embedded mechanism for sensing the size and orientation of each pinch, we are motivated to explore whether fabrics such as those reported in [9] and [6] would be usable for such an interface, or what modifications would need to be made for practical use. Accurately sensing the rich interactions possible with fabric, such as folding, twisting and sliding, is a challenge. For example, grabbing a piece of fabric can result in a wide variety of fabric effects (Figure 4), with different amounts of material manipulated, that can create a rich form of input, or else sometimes make it exceedingly difficult to ascertain the user’s intent. Our proposed interface needs to accurately detect the angle of the initial grab. Although [6] used a simple PCA algorithm for that purpose, it was not reliable for detecting a fold at an arbitrary axis. Ideally, this would be possible up to the limit of human performance, which the earlier experiments demonstrated lies somewhere in the 30° to 45° range.



Figure 4: A variety of fabric effects from a simple pinch. Note the differences in twisting, and puckering of the fabric with different resulting numbers of folds.

4.1 Grabrics implementation

Grabrics is a sensor that is, except for the sensing microcontroller, entirely made of textile materials. It consists of a 90×90 mm sensing surface made of 30 hexagonal conductive pads with 10 mm diameter. The pads are made of conductive thread (Shieldex[®] 235/34) embroidered in a circular pattern to avoid a directional preference when two of these pads slide over each other. The user

interacts with the sensor by grabbing a fold into the fabric, which results in an interconnection between some of the pads. A microcontroller board (Tiva C Series ARM Cortex-M4) connected directly to the fabric senses these interconnections at an update rate of 6.25 Hz and reports that information to a host computer connected via a USB cable for further analysis.

Before describing the experiment and results, it is first important to understand in greater detail the capabilities and limitations of Grabrics. In the current implementation, the Grabrics electronics can measure each pad, one at a time, and determine what other pads are touching it either directly or through other intervening pads. These form a single blob, and it is important to note that it is not possible to know which specific pads are connected to one another within this group. However, there can be separate sets of activated pads (blobs) that are electrically disconnected from one another, which the existing electronics can indeed detect. We hypothesized that this information may be important for distinguishing the direction of a given fold.

4.2 Experiment setup

We recruited 9 participants (7 male, 2 female, ages 22-31, all right handed) from the university community. Two of the male participants had also completed the first experiment. Since this was a test of the capabilities of the Grabrics fabric itself when coupled with a detection algorithm, the fabric was placed on a table in front of the participant. This removed any issues with situating the electronics on their arm, which could interfere with their grabbing since in the current prototype the electronics are fairly bulky and somewhat fragile. Unlike in the first experiment, the user was allowed to look at the fabric while grabbing. This precluded any error due to the participants’ proprioception or misjudgement of the location and orientation of the fabric, ensuring that the experiment is measuring the performance of the prototype and the algorithm as much as possible while still using realistic grasping motions. We expect that the types of grabs will be similar to those when the fabric is worn, although some adjustments may need to be made.




Figure 5: A microcontroller senses the connections between the conductive pads on the fabric and transmits that information to a host computer for further analysis.

After signing a consent form, participants were shown the prototype and allowed to familiarize themselves with it by folding it and watching a real-time visualization on a laptop screen in front of them that let them see how the prototype functioned. The experimenter ensured they correctly understood target angles drawn on a sheet of paper. Once they understood the angles, the experi-

menter showed the participant a piece of paper with the target angle to pinch. The participant then pinched the fabric and held it for approximately five seconds to ensure a stable reading from the sensors. When the experimenter saw that the pinch was stable, she pushed a button on the logging computer that captured the Grabrics sensor data at that instant. Participants were allowed to smooth the Grabrics if they desired after each trial, as well as adjust its position slightly on the table to a comfortable position. However, the experimenter ensured that its angle on the table was consistently upright in front of the participant throughout the experiment. Eight times during the experiment participants indicated they had pinched the wrong angle immediately after they had done so. They were then allowed to perform the pinch a second time with the proper orientation, and the mistaken angles were removed from the data set.

Four pinch angles were tested ($0^\circ - 135^\circ$ in 45° increments) in random order, with each repeated 20 times by participant 1, and 30 times by the other eight participants. The zero axis was vertical on the table in front of the user, with angles increasing clockwise. In addition, participants were asked to do a fifth "random" pinch to gather additional data on other patterns that could arise in practice. However, these were not used in the analysis, but rather just for gathering anecdotal information on other possible fabric folding patterns. After every 50 trials, participants were given a break to prevent fatigue. The experiment lasted approximately 20 minutes per participant.

4.3 Results

With the labeled data from the experiment, consisting of 1040 total points, we could evaluate different methods for computing the angle of the fold.

The median number of active pads during a fold at any of the tested angles was 10 pads ($SD = 2.5$), which is about a third of the total Grabrics pads. Due to broken connections at the interface between the fabric and the Grabrics PCB, three pads were not functional during the experiment (Figure 6), and were excluded from analysis. We note that all three are at the periphery of the active area, resulting in a somewhat smaller total active surface than if all were functioning properly.




Figure 6: The Grabrics sensor textile with the 30 conductive pads and the connection clip for the microcontroller on the right. Due to mechanical issues, three pads at the periphery of the Grabrics active area were disconnected (highlighted in red).

Our first approach was to try and generate explicit rules for determining the fold angles, based on algorithms such as principle component analysis (PCA) or clustering techniques. These were based

on the observation that in some cases, it is relatively unambiguous where the fold lies. In the simplest case, with a straight line of pads active, they are clearly folded on top of each-other, and thus the fold must run in a line roughly perpendicular to a line drawn between them (Figure 7.a). However, in other cases, given the limitations of sensing in the Grabrics prototype, it is impossible to tell for certain which direction the fabric is folded, such as when four pads form a single electrical blob in a non-linear arrangement (Figure 7.b). There are many possible cases between these two extremes, with a wide variety of pads activated in different groups. Although we designed several explicit algorithms using clustering to determine the fold angle, these approaches tended to only achieve approximately 50% accuracy rates. We suspect that this is at least partially due to the fact that in ambiguous cases, these clustering algorithms do not take into account the fact that some positions and angles for a given set of connected pads are more likely simply due to human physiology and how the hand rotates and grabs. Such techniques could be improved, but initial results using more general machine learning approaches, which do take these factors into account, were more encouraging, so we focused on them instead.




Figure 7: In (a) we see four pads in a row activated. Only one fold, perpendicular to the line of pads, can possibly generate this pattern. In (b) also with all four pads connected to each-other, at least two different viable folds (vertical and horizontal) will generate this same pattern, so we cannot know for certain which is correct.

The caret package [19] in R [25] was used to generate a predictive model, using random forest (method = "rf") with 10 repetitions of 10-fold cross-validation (method = "repeatedcv", number = 10, repeats = 10). Two different analyses were carried out. First, a simple vector representing each pad as either connected or unconnected was used, with no information about which blobs are electrically isolated. This achieved a classification accuracy of 81% ($\kappa = 0.75$, $accuracySD = 0.03$). Adding information about which pads were electrically connected to which others, which takes into account electrically isolated blobs of connected pads, slightly improved the accuracy to 85% ($\kappa = 0.80$, $accuracySD = 0.03$). The confusion matrix for the latter final model is provided in Table 3.

		Predicted				class error
		0°	45°	90°	135°	
Observed	0°	229	9	14	8	12%
	45°	8	222	13	18	15%
	90°	30	13	212	4	18%
	135°	14	18	7	221	15%

Table 3: Confusion matrix for random forest classification of pinched angle based on all information available from the prototype for electrical connections between pads. Overall accuracy 85%.

It is worth noting that accuracy of the final model increases to 91% ($\kappa = 0.75$, accuracySD = 0.03) when restricting the analysis to only 0° and 90° angles.

4.4 Discussion

The Grabrics sensor has two physical limitations: First, the conductive pads and connecting wires are large and relatively stiff, which makes it physically difficult for the material to easily conform to a fold along arbitrary axes, simply due to the limited way that the fabric “naturally” creases. Second, even if it were physically possible to accurately fold the fabric at arbitrary angles, the pads are relatively large, such that the resolution of the information obtained by the electronics may not be enough to be sure of the intended angle. For example, given the large pad size, it is impossible to ascertain the precise angle of two connected pads since they can be slid side-to-side against each-other without losing the connection. This means the angle of the fold can change without any corresponding change in the detected pads. As already noted, this limits the resolution of the detection.

Especially given that we expect future implementations to improve on these issues, we investigated methods for detecting the orientation of the fold that would work in real-time, with the assumption that if it is possible on the existing coarse prototype, then it will only improve with further prototype refinements.

The results show that even with this issue, a promising classification accuracy is achieved, but it would need to be improved to make it reliable enough for real-world use. Although we could explore improvements to the classification, we expect that upgrading the Grabrics prototype would be the more fruitful approach.

In this second experiment, participants reported that pinching Grabrics was challenging when the pads were aligned at the pinch axis, and due to the stiffness of the pads it was harder to create the desired fold reliably. Since the pads are stiffer than the surrounding fabric, folds that mostly crease fabric between the pads would likely be the easiest to perform. With the current pad configuration, this means (with the electronics to the left) along the vertical axis, and at 60° off vertical in either direction (Figure 8). However, we expect that this effect is less apparent when doing larger grabs, as participants were instructed to do in the experiment described in this section. This is because the larger the pinch, the less necessary it is to fold the pads themselves, since the fabric can more easily conform around the stiff pads in the looser fabric at the centre of the grab.

While this could be perceived as a limitation, one could also take advantage of textile characteristics to enforce or facilitate certain pinch angles. For example, a marking menu interface may benefit from a discrete number of possible angles, as constrained by the fabric, since it may subtly encourage the user to fold the fabric in a way that makes the fold angle easier to assess. These issues also imply that frequently accessed or more important menu items should be mapped to axes with higher performance. In summary, for best performance, the character of the fabric may need to be tuned to the intended application, and mappings of menu items to specific angles may need to take into account which angles are physically easier to perform or to detect, to minimize errors.

5. CONCLUSIONS & FUTURE WORK

We have reported on two experiments that approach the problem of fabric interfaces similar to marking menus from two perspectives. First, we demonstrated that the limit of human performance in grabbing common fabrics is in the range of $30^\circ - 45^\circ$ angles. Second, we showed that even with a relatively coarse prototype, we can achieve classification accuracy, at 45° increments, of 85%




Figure 8: The fold angles that best avoid folding the pads themselves are along the vertical, as well as 60° off the vertical in either direction.

when using a straightforward machine learning algorithm such as random forests. When taking these two results together, we contend that a practical interface using such a fabric to select and manipulate different values will soon be achievable. In fact, with proper feedback to mitigate the impact of classification errors, even the current system may be usable, especially if restricted to only two fold angles (0° and 90°), which achieved over 90% accuracy.

Our prototype only had a coarse resolution due to technical limitations of the production process. For future implementations, making the pads smaller and/or more flexible will allow the fabric to feel more supple and will allow it to be easily pinched in different axes. We expect that smaller conductive pads, resulting in higher resolution, should also allow finer angle and rolling control than the existing Grabrics prototype. Although as previously noted, some applications may benefit from carefully designed fabrics that constrain folds to certain axes, we expect that it would still be beneficial to have higher resolution detection of the fold axis.

We also propose examining the classification confusion matrix (Table 3) in greater detail to determine what might account for the (for example) greater error rate in distinguishing 0° vs. 90° folds as opposed to 90° vs. 135° . At first glance, one would expect that the further apart the angles, the easier they should be to distinguish. This is not necessarily the case, potentially due to the way Grabrics folds along different axes, and understanding exactly why may lead to improved designs. In addition, once we put Grabrics on the arm itself, certain angles may be easier to pinch, e.g. along the axis of the arm, since the fabric will already be conforming around this axis due to the curvature of the arm. By aligning the pads in different orientations on the body, we may be able to take advantage of such effects to reinforce the bias caused by the stiffer pads in the fabric itself, and perhaps improve the classification rate.

In the first study we found that fabric had no influence on users’ pinching accuracy or comfort, although one stiffer fabric was excluded since it was not possible to pinch it. Thus, although we know that some fabric types will not work for such an interface, if a fabric is able to be pinched, it appears that performance is equivalent across a range of fabrics. Further investigations are needed to explore the effect of other factors such as arm position, the user’s mobility, or feedback (auditory or haptic) as investigated by [2]. In particular, both of our studies were conducted in a laboratory environment. When distracted in the real-world, not only the user’s capabilities are important, but the Grabrics sensor itself needs to be robust enough to be used while moving. Furthermore, this paper

only looked at grabbing a fold in the fabric. We reported that fabric type influenced users perceived easiness of rolling, but additional investigations are needed to evaluate fabric rolling and characterize any factors that might, for example, reduce rolling granularity or range. Last, we focused on using Grabrics on the forearm, and our results may not hold for other areas of the body that would be potentially useful, such as the upper leg.

Although marking menus have been well-studied, an implementation with a wearable fabric controller may have implications for their usability. This would need to be tested by implementing marking menus and performing a user study to assess how learnable such a system would be in practice.

6. ACKNOWLEDGMENTS

We would like to thank the participants of our study. This project was funded by the German Federal Ministry of Education and Research under the project “Intuitex” (Grant 16SV6264K), as well as a McGill Engineering Doctoral Award, CIRMMT Inter-Centre Research Exchange Funding, and an NSERC PGS-D Scholarship. McGill REB #446-0416.

7. REFERENCES

- [1] J. Bauer, A. Ebert, O. Kreylos, and B. Hamann. Marking menus for eyes-free interaction using smart phones and tablets. In *Availability, Reliability, and Security in Information Systems and HCI*, pages 481–494. Springer, 2013.
- [2] S. Brewster, J. Lumsden, M. Bell, M. Hall, and S. Tasker. Multimodal ‘eyes-free’ interaction techniques for wearable devices. In *CHI '03*. ACM, 2003.
- [3] N. Chen, F. Guimbretiere, M. Dixon, C. Lewis, and M. Agrawala. Navigation techniques for dual-display e-book readers. In *CHI '08*. ACM, 2008.
- [4] D. Dobbstein, P. Hock, and E. Rukzio. Belt: An Unobtrusive Touch Input Device for Head-worn Displays. In *CHI '15*. ACM, 2015.
- [5] G. Gioberto, J. Coughlin, K. Bibeau, and L. E. Dunne. Detecting bends and fabric folds using stitched sensors. In *ISWC '13*. ACM, Sept. 2013.
- [6] N. Hamdan, F. Heller, C. Wacharamanotham, J. Thar, and J. Borchers. Grabrics: A foldable two-dimensional textile input controller. In *CHI '16 EA*. ACM, 2016.
- [7] S. G. Hart. Nasa-task load index (nasa-tlx); 20 years later. *Proc. of the human factors and ergonomics society annual meeting*, 50(9):904–908, 2006.
- [8] F. Heller, S. Ivanov, C. Wacharamanotham, and J. Borchers. FabriTouch: Exploring Flexible Touch Input on Textiles. In *ISWC '14*. ACM, 2014.
- [9] F. Heller, H.-Y. K. Lee, P. Brauner, T. Gries, M. Ziefle, and J. Borchers. An intuitive textile input controller. In *Mensch und Computer 2015 – Proceedings*. De Gruyter Oldenbourg, 2015.
- [10] G. Herkenrath, T. Karrer, and J. Borchers. Twend: twisting and bending as new interaction gesture in mobile devices. In *CHI '08 EA*. ACM, 2008.
- [11] P. Holleis, A. Schmidt, S. Paasovaara, A. Puikkonen, and J. Häkkinen. Evaluating Capacitive Touch Input on Clothes. In *MobileHCI '08*. ACM, Sept. 2008.
- [12] T. Karrer, M. Wittenhagen, L. Lichtschlag, F. Heller, and J. Borchers. Pinstripe: eyes-free continuous input on interactive clothing. In *CHI '11*. ACM, May 2011.
- [13] M. Khalilbeigi, R. Lissermann, W. Kleine, and J. Steimle. Foldme: interacting with double-sided foldable displays. In *TEI '12*. ACM, 2012.
- [14] G. Kurtenbach and W. Buxton. The limits of expert performance using hierarchic marking menus. In *CHI '93*. ACM, 1993.
- [15] S.-S. Lee, S. Kim, B. Jin, E. Choi, B. Kim, X. Jia, D. Kim, and K.-P. Lee. How Users Manipulate Deformable Displays As Input Devices. In *CHI '10*. ACM, 2010.
- [16] J. Lepinski and R. Vertegaal. Cloth displays: Interacting with drapable textile screens. In *TEI '11*. ACM, 2011.
- [17] R. Lissermann, J. Huber, A. Hadjakos, S. Nanayakkara, and M. Mühlhäuser. EarPut: Augmenting Ear-worn Devices for Ear-based Interaction. In *OZCHI '14*. ACM, 2014.
- [18] G. N. Marentakis and S. A. Brewster. Effects of feedback, mobility and index of difficulty on deictic spatial audio target acquisition in the horizontal plane. In *CHI '06*. ACM, 2006.
- [19] Max Kuhn. Contributions from Jed Wing and Steve Weston and Andre Williams and Chris Keefer and Allan Engelhardt and Tony Cooper and Zachary Mayer and Brenton Kenkel and the R Core Team and Michael Benesty and Reynald Lescarbeau and Andrew Ziem and Luca Scrucca and Yuan Tang and Can Candan. *caret: Classification and Regression Training*, 2016. R package version 6.0-68.
- [20] I. Oakley and J. Park. Motion marking menus: An eyes-free approach to motion input for handheld devices. *Int. Journal of Human-Computer Studies*, 67(6):515–532, 2009.
- [21] H. Perner-Wilson, L. Buechley, and M. Satomi. Handcrafting textile interfaces from a kit-of-no-parts. In *TEI '11*. ACM, 2011.
- [22] J. Peschke, F. Göbel, T. Gründer, M. Keck, D. Kammer, and R. Groh. Depthtouch: An elastic surface for tangible computing. In *AVI '12*. ACM, 2012.
- [23] S. Pook, E. Lecolinet, G. Vaysseix, and E. Barillot. Control menus: Execution and control in a single interactor. In *CHI '00 EA*. ACM, 2000.
- [24] H. P. Profita, J. Clawson, S. Gilliland, C. Zeagler, T. Starner, J. Budd, and E. Y.-L. Do. Don’t Mind Me Touching My Wrist: A Case Study of Interacting with On-body Technology in Public. In *ISWC '13*. ACM, 2013.
- [25] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. ISBN 3-900051-07-0.
- [26] J. Thar. Pinstripe: Integration & evaluation of a wearable linear input controller for everyday clothing. Bachelor’s thesis, RWTH Aachen University, February 2013.
- [27] B. Thomas, K. Grimmer, J. Zucco, and S. Milanese. Where Does the Mouse Go? An Investigation into the Placement of a Body-Attached TouchPad Mouse for Wearable Computers. *Pers. and Ubiqu. comp.*, 6(2):97–112, Jan. 2002.
- [28] G. M. Troiano, E. W. Pedersen, and K. Hornbæk. User-defined gestures for elastic, deformable displays. In *AVI '14*. ACM, 2014.
- [29] M. Weigel, T. Lu, G. Bailly, A. Oulasvirta, C. Majidi, and J. Steimle. iSkin: Flexible, Stretchable and Visually Customizable On-Body Touch Sensors for Mobile Computing. In *CHI '15*. ACM, 2015.
- [30] S. Zhao, P. Dragicevic, M. Chignell, R. Balakrishnan, and P. Baudisch. Earpod: Eyes-free menu selection using touch input and reactive audio feedback. In *CHI '07*. ACM, 2007.