

What Caused that Touch? Expressive Interaction with a Surface through Fiduciary-Tagged Gloves

Nicolai Marquardt, Johannes Kiemer, Saul Greenberg

Department of Computer Science
University of Calgary, 2500 University Drive NW, Calgary, AB, T2N 1N4, Canada
[nicolai.marquardt, jkiemer, saul.greenberg]@ucalgary.ca

ABSTRACT

The hand has incredible potential as an expressive input device. Yet most touch technologies imprecisely recognize limited hand parts (if at all), usually by inferring the hand part from the touch shapes. We introduce the fiduciary-tagged glove as a reliable, inexpensive, and very expressive way to gather input about: (a) many parts of a hand (fingertips, knuckles, palms, sides, backs of the hand), and (b) to discriminate between one person's or multiple peoples' hands. Examples illustrate the interaction power gained by being able to identify and exploit these various hand parts.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors

Keywords: Surfaces, touch, gestures, postures, gloves, fiduciary tags, multi user

INTRODUCTION AND RELATED WORK

The surge in availability of touch technologies for both small devices and large surfaces has introduced direct touch to the public as a popular and powerful way to interact with computers. Earlier technologies handled only a single touch, where touch often emulated a mouse to allow a person to interact with a traditional GUI. Later technologies provided two-touch and then multi-touch. This favoured the development of completely new touch and gestural interaction techniques not afforded by a mouse.

Most commercial technologies do not discriminate what caused the touch. At best, some assume that people tend to touch and gesture with their hands: single touch by fingertip, two-touch by two fingertips, and multi-touch by multiple fingers and whole-hand postures (e.g., [14]).

Our own research asks: what body part causes a touch, and is it useful to know that? We assert that more interaction power can be expressed if the system knows what hand part is causing the touch. We are not alone in this assertion: others have developed a variety of algorithms and technologies – some for domains other than surface interaction – to recognize particular fingers, or hand postures, or to distinguish between

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ITS'10, November 7–10, 2010, Saarbrücken, Germany.

Copyright © 2010 ACM 978-1-4503-0399-6/10/11...\$10.00



Figure 1. The fiduciary-tagged glove (top, bottom, side).

the hands of one person or between multiple people. Perhaps the most common method to recognize bare hands uses vision. Kruger's groundbreaking VideoPlace [7] tracked silhouettes, and used that to determine hand postures. Newer techniques recognize bare fingers (e.g., [8]), and can even track translation and rotation parameters associated with particular bare hand gestures (e.g., [11]). The Microsoft Surface visually identifies fingertip blobs, where it uses that shape to return the fingertip position and orientation. Similarly, Dang [3] uses blobs to distinguish if finger touches are from the same or different hands, while Freeman [5] offers further distinction of hand postures via vision techniques. Specialized hardware also affords abilities to distinguish hand parts. Examples include the MERL DiamondTouch surface that identifies which touch came from which person [4] and finger and posture recognition inferred by the shapes returned from it [16], Benko's EMG muscle sensing approach that identifies fingers touching a surface [1], and fingerprint recognition devices [6].

Our contributions build upon these attempts. In particular, our method knows what hand parts are touching the surface, and which / whose hand the touches are attached to. Importantly, it differs from prior work as we offer a very inexpensive, simple, yet highly accurate method: gloves tagged with fiduciary tags (Figure 1). We show how expressive interactions can be created given the certainty of this knowledge. While this approach requires people to wear gloves, it allows interaction designers to rapidly prototype and explore meaningful interfaces. This can occur well before research solves the problem of accurate bare-hand touch identification.

Glove-based tracking is well known in other domains. Within virtual reality environments, [12] surveys a multitude of ways to track gloves: optical tracking, markers, silhouette analysis, magnetic tracking, optical fibres, while [13] use coloured gloves to discriminate parts of the hand.

Within augmented reality, fiduciary tags attached to gloves were used to track gestures made in 3D-space, and as a way for a person to handle virtual objects in that space (e.g., [2]). To our knowledge, no one has used fiduciary-tagged gloves *on surfaces*, nor has anyone investigated the surface interaction possibilities afforded by them.

THE FIDUCIARY-TAGGED GLOVE

The Microsoft Surface is capable of recognizing multiple fiduciary tags placed atop of it [9]. Through the Surface API, the programmer can easily retrieve information associated with the tag: its 8-bit identification number, its coordinates of where it touches the surface, and its orientation to the surface normal. We exploit this Surface capability.




Figure 2. A fiduciary-tagged glove touching the tabletop surface.

Figure 1 illustrates our fiduciary-tagged glove. Figure 2 shows it touching a Surface application that visualizes the id, coordinates, and orientation information returned from a recognized tag. There is nothing special about our glove, except that it is thin (for comfort), and was cheaply available at a local clothing store. The only constraint is that its material is non-reflective of infrared (IR) light. We stuck fiduciary tags onto it to allow 15 key hand parts to be identified (Table 1). We use 2x2 cm fiduciaries, which are currently the smallest size recognized reliably by the Microsoft Surface. Because the Microsoft Surface only recognizes flat tags, we could not simply print the tag onto the glove fabric itself. In our current version (pictured in Figure 1) we glue small squares of plastic into the inside of the glove, where they lie underneath the part of the glove to be tagged. We then stick the fiduciary onto the glove itself, atop these flat surfaces (we originally had the plastic on the outside, but this leaves exposed edges which catch on things). The result is a reasonably comfortable and very lightweight glove where edges do not catch. Costs are minimal: our glove was purchased for under \$10, and label paper for a few dollars. All other materials are readily available around the home or workplace.

What is important is that these small tags can identify the major parts of a hand. The version shown in Figure 1 has 15 tags, which means that it can distinguish between the 15 hand parts listed in Table 1. Furthermore, if a person is wearing two gloves, or if more than one person is wearing gloves, then we can also distinguish which hand and/or which person that tag is attached to.

CALIBRATING THE GLOVE

Once the tags are attached to a new glove, a one-time calibration is necessary for identifying: the particular area of the hand that correspond to a fiduciary tag, left vs. right hands for a person wearing two gloves, and the identity of the person to discriminate between multiple people wearing gloves. The entire process takes only a few minutes.

Table 1. Recognized hand parts, and the # of tags required

each fingertip (5)	palm (1)
back of each finger (5)	wrist (1)
back of the hand (1)	side of the hand (2)

Our system guides a person through a simple calibration process. It displays life-sized top, bottom and side views of a flat hand (Figure 3 left), where the images outline the 15 hand areas (Table 1) that can be associated with the tags. First, the person specifies which hand is being calibrated; the view changes accordingly to show the left or right hand. Second, the person optionally identifies themselves by name, which allows those tags to be associated – and thus identify – a particular person. Third, the person associates their fiduciary tags to corresponding hand parts by placing their gloved hand on top of a particular view. The system recognizes the tags that touch particular image areas, and will store that as an association. This process is repeated for the back of the hand and the side. Visual feedback (Figure 3 right) indicates the recognized fiduciary tag. The system also checks to see if the same fiduciary tag is incorrectly assigned to two or more areas of the hand. Corrections are easily done, where the person can just touch the particular image area with the corresponding tag. The configuration of each glove (and the optional user name) is persistently saved to an XML file – making it easy to access this information from within each developed application on the touch table. The calibration process can then be repeated for additional gloves and/or people.




Figure 3. Glove calibration process.

DIMENSIONS OF RECOGNIZED INPUT

The following three input dimensions are available through tracking the fiduciary-tagged glove: identifying fingers and other parts of the hand; postures; and hand gestures.

Identifying individual parts of the hand that are touching the surface is straightforward and highly reliable. It merely requires looking up the hand part associated with the returned tag identity. Thus particular finger tips can be easily distinguished, as well as the other hand parts listed in Table 1. As each tag provides the accurate pixel position and its orientation in degrees, we also know the same information for that hand part. Because multiple tags can be tracked simultaneously, we also know what hand parts are concurrently touching the surface. We also know what hand and what person is causing that touch, because hand and person identity are associated with tags.

Identifying hand postures atop the surface is done by monitoring the relative *distance* and *orientation* between two or more particular tags. Our system defines a *posture*




Figure 4. Identify postures through changed tag orientation.

as a particular subset of glove tags that fall between a given range of relative positions and orientations. When a person lays their hand on the table, the system sees the multiple individual tags, then sees if their identity, relative position and orientation matches its posture set, and if so, returns that posture.

Two approaches illustrate posture recognition. The first is *orientation angle*. Consider two tags attached to the side of the hand (Figure 4). When the side of the hand is in a straight orientation, the angle between the two tags is close to 0° (Figure 4a). The angle changes to around 90° with an L-shaped posture, to 120° with a curved C-shaped posture, and around 180° with a fist (as shown in Figures 4b-d). Thus the sole comparison of the relative *orientation angle* between these two particular tags at the side of the hand allows reliable identification of four different postures.

Our second approach interprets the *relative distance* (instead of orientation) between specific tags to differentiate between hand postures. Consider, for example, tags attached to the fingertips, palm, and wrist of the glove around the user's hand. If the hand is pressed flat against the surface all of these tags are detected simultaneously. When spreading the fingers the relative distance between the fingers increase, and when bending the fingers – like when grabbing an object – the distance between the fingers and the palm decrease, and thus these postures can be discriminated.

Transitions between postures are also recognized. As a person changes postures while touching the surface, the changes in relative distance and orientation of the tags are also continuous. Thus we can easily track the fluent transitions between these postures. This capability is also essential for gesture recognition.

Identifying gestures is done by tracking the changes in the position and orientation of each tag over time. Gesture recognition algorithms (e.g., [15]) can thus observe gestures performed by the user. Gesture recognition can be applied to individual tags (e.g., circle gesture of a single finger), and to groups of two or more tags (e.g., pinching with two fingers). Importantly, the unique identification of each part of the hand allows limiting the recognition of certain gestures to only particular fingers or parts of the hand. For instance, the pinch-to-zoom gesture can be limited to index finger and thumb, or a swipe gesture is recognized only if performed with the side of the hand. A gesture that appears similar (e.g., a pinch gesture done by index features of two hands, or by similar movement of the finger of two different people) can be assigned completely different meanings (or ignored as noise), if desired.




Figure 5. Interactive drawing application on the digital tabletop, using the identity of touches for interaction.

EXPRESSIVE INTERACTION

Of course, much prior work involves identifying fingers, postures and gestures. What makes our fiduciary-tagged glove unique is that this identification is easily done, highly reliable, and – importantly – that the meaning of the touches, postures and gestures can be unique to the particular hand parts, hands, and people involved. This reliability and added meaning leads to what we call *expressive interaction*.

We now illustrate how such touches, postures, and gestures can be applied to expressive interaction of a single person using one or two hands, and of multiple people interacting simultaneously. Similar to [1], we use a deliberately simple drawing application on a digital surface. We see this as a “Hello World” application in the domain of touch identification. Through its simplicity, it allows us to focus on illustrating the expressive forms of interaction. Other, more advanced, applications are possible and part of future work.

One person, one handed interaction.

A single person interacting with the digital surface can have unique functions assigned to each of the marked areas of the glove. In our drawing application, strokes are painted by touching the finger tips to the surface, where each finger can uniquely define different stroke properties. For example, a person can use the system interface to dynamically assign a different stroke width, color, and opacity to each finger (Figure 5, left). Other parts of the hand can have different drawing functions. Panning is done by pressing the flat hand against the canvas and moving it around, and saving the large drawing area is done by a swiping gesture with the back of the hand. Knuckles erase regions as swaths, where the eraser size depends on the knuckle used.

One person, two handed interaction.

When hand posture and gesture is included, the system has to determine whether particular touches and touch movements represent a given hand posture or gesture, or just arise from the accidental arrangements of other tags touching the surface. These cases are easily distinguished with the fiduciary-tagged glove, because the particular touch points are identified, via the tags, with particular hand parts and with hands. Thus the system can correctly interpret touch points and their assignment to particular hand postures and gestures, even when touches from other hand parts resemble particular posture or gesture configurations.

As with individual tags, particular hand postures (through the identification of two or more tags) – even those that are quite similar in appearance – can be assigned to specific drawing

functions in our example application. For instance, a region of the canvas can be selected and copied by defining a rectangle area on the canvas by forming two L-shaped selector tools with the thumb and index finger of both hands – the area in between the two L-shaped postures is the selected area (see Figure 5 right). To paste this copied area to the other positions on the drawing canvas, the user can simply form the same L-shape, but this time using the thumb and pinkie finger. Other postures for copy/paste are feasible, but the major point is that with fiduciary gloves we can easily differentiate postures that are almost identical from a computer vision perspective.

Two people interacting simultaneously.

In the simplest case, two (or more) people interact independently on individual tasks on the same interactive tabletop surface. Because the system knows about the association of each contact point to an individual person, there is no interference between all the simultaneous touches. For example, if two people simultaneously produce a single touch, this cannot be mistakenly recognized as a two-touch posture or gesture (such as a two-finger-pinch gesture) of one person.

The information about identity can also be leveraged to only allow certain actions to be performed by a particular person. For instance, objects created by one person can only be deleted by the same person; or the access to certain areas on the surface or displayed widgets (like buttons, sliders) can be restricted or tailored to a particular user (e.g., [4]). If consensus is required, the system may require all current users to perform a combined posture or gesture simultaneously (e.g., *multi-user coordination policies* in [10]). For example, to switch to a new drawing, which could have a large impact on the workspace, all participants may have to place their fists on the surface.

Interaction with tagged tangible objects on the surface

Fiduciary tags were originally envisioned as a way for the surface to recognize the position and orientation of tagged physical objects placed atop of it [9]. We combine such tagged objects to interact with particular hand parts identified by our gloves.



Figure 6. Calibrating a finger's color with a tangible object.

Our drawing application uses this capability to allow a person to rapidly reconfigure the color capabilities of each fingertip (Figure 6). We use a (3x3x3 cm) tagged physical cube, whose side, position and orientation is detected by the surface. In this case, three sides represent controls for adjusting hue, saturation and luminance. When a person places a finger tip close to that cube, the surface visually indicates that the two are connected (the two connected circles in Figure 6, which also shows that finger's current drawing color). Depending on the side of the cube placed on the surface, the person then rotates the cube to dynamically reconfigure the color hue, saturation or luminance associated with that finger.

CONCLUSION

Our fiduciary-tagged glove introduces an inexpensive but expressive way to gather accurate input about many parts of a hand (e.g., the 15 hand parts in Table 1), what hand it is from, and owner identification. The comparison of relative orientation and distance between particular tags facilitates posture and gesture recognition, even when postures and gestures from different hand parts and hands appear similar. Expressive interaction occurs when the meaning of the touches, postures and gestures can be unique to the particular hand parts, hands, and people involved.

While deceptively simple, fiduciary-tagged gloves provide a cheap yet effective means for others to explore expressive interactions. We believe that exploiting the identity of hand parts can lead to interaction techniques that go far beyond current methods that cannot differentiate touches reliably. We have just touched the surface of what is possible.

ACKNOWLEDGMENTS Partially funded by the iCORE/NSERC/SMART Chair in Interactive Technologies, Alberta Innovates, NSERC, and SMART Technologies Inc.

REFERENCES

1. Benko, H., Saponas, T.S., Morris, D., and Tan, D. Enhancing input on and above the interactive surface with muscle sensing. *Proc. of ITS*, ACM (2009), 93-100.
2. Buchmann, V., Violich, S., Billingham, M., and Cockburn, A. FingARTips: gesture based direct manipulation in Augmented Reality. *Proc. of GRAPHITE*, ACM (2004).
3. Dang, C.T., Straub, M., and André, E. Hand distinction for multi-touch tabletop interaction. *Proc. of ITS*, ACM (2009).
4. Dietz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. *Proc. of UIST*, ACM (2001), 219-226.
5. Freeman, D., Benko, H., Morris, M.R., and Wigdor, D. ShadowGuides: visualizations for in-situ learning of multi-touch and whole-hand gestures. *Proc. of ITS*, ACM (2009).
6. Holz, C. and Baudisch, P. The generalized perceived input point model and how to double touch accuracy by extracting fingerprints. *Proc. of CHI*, ACM (2010), 581-590.
7. Krueger, M.W., et al. VIDEOPLACE - An Artificial Reality. *SIGCHI Bull.* 16, 4 (1985), 35-40.
8. Letessier, J. and Bérard, F. Visual tracking of bare fingers for interactive surfaces. *Proc. of UIST*, ACM (2004).
9. Microsoft Inc. Tagged Objects. MSDN Surface SDK 1.0, 2010.
10. Morris, M.R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Beyond "social protocols": multi-user coordination policies for co-located groupware. *Proc. of CSCW*, ACM (2004).
11. Schlattman, M. and Klein, R. Simultaneous 4 gestures 6 DOF real-time two-hand tracking without any markers. *Proc. of VRST*, ACM (2007), 39-42.
12. Sturman, D.J. and Zeltzer, D. A Survey of Glove-based Input. *IEEE Comput. Graph. Appl.* 14, 1 (1994), 30-39.
13. Wang, R.Y. and Popović, J. Real-time hand-tracking with a color glove. *Proc. of SIGGRAPH*, ACM (2009), 1-8.
14. Wobbrock, J.O., Morris, M.R., and Wilson, A.D. User-defined gestures for surface computing. *Proc. of CHI*, ACM (2009).
15. Wobbrock, J.O., Wilson, A.D., and Li, Y. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. *Proc. of UIST*, ACM (2007), 159-168.
16. Wu, M. and Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *Proc. of UIST*, ACM (2003), 193-202.