

Extended Multitouch: Recovering Touch Posture and Differentiating Users using a Depth Camera

Sundar Murugappan¹, Vinayak¹, Niklas Elmqvist², Karthik Ramani^{1,2}

¹School of Mechanical Engineering and ²School of Electrical and Computer Engineering
Purdue University, West Lafayette, IN 47907, USA
{smurugap,fvinayak,elm,ramani}@purdue.edu




Figure 1. Extracting finger and hand posture from a Kinect depth camera (left) and integrating with a pen+touch sketch interface (right).

ABSTRACT

Multitouch surfaces are becoming prevalent, but most existing technologies are only capable of detecting the user's actual points of contact on the surface and not the identity, posture, and handedness of the user. In this paper, we define the concept of *extended multitouch interaction* as a richer input modality that includes all of this information. We further present a practical solution to achieve this on tabletop displays based on mounting a single commodity depth camera above a horizontal surface. This will enable us to not only detect when the surface is being touched, but also recover the user's exact finger and hand posture, as well as distinguish between different users and their handedness. We validate our approach using two user studies, and deploy the technique in a scratchpad tool and in a pen + touch sketch tool.

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Interaction styles*; I.3.6 Computer Graphics: Methodology and Techniques—*Interaction techniques*

General Terms

Design, Algorithms, Human Factors

Author Keywords

Multitouch, tabletop, depth camera, pen + touch, evaluation.

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 publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UIST'12, October 7–10, 2012, Cambridge, Massachusetts, USA.

Copyright 2012 ACM 978-1-4503-1580-7/12/10...\$15.00.

INTRODUCTION

Multitouch surfaces are quickly becoming ubiquitous: from wristwatch-sized music players and pocket-sized smartphones to tablets, digital tabletops, and wall-sized displays, virtually every surface in our everyday surroundings may soon come to life with digital imagery and natural touch input. However, achieving this vision of ubiquitous [27] surface computing requires overcoming several technological hurdles, key among them being how to augment any physical surface with touch sensing. Conventional multitouch technology relies on either capacitive sensors embedded in the surface itself, or rear-mounted cameras capable of detecting the touch points of the user's hand on the surface [10]. Both approaches depend on heavily instrumenting the surface, which is not feasible for widespread deployment.

In this context, Wilson's idea [30] of utilizing a front-mounted depth camera to sense touch input on any physical surface is particularly promising for the following reasons:

- **Minimal instrumentation:** Because they are front-mounted, depth cameras integrate well with both projected imagery as well as conventional screens;
- **Surface characteristics:** Any surface, digital or inanimate, flat or non-flat, horizontal or vertical, can be instrumented to support touch interaction;
- **On and above:** Interaction both on as well as above the surface can be detected and utilized for input;
- **Low cost:** Depth cameras have recently become commodity products with the rise of the Microsoft Kinect, originally developed for full-body interaction with the Xbox game console, but usable with any computer.

In this paper, we define a richer touch interaction modality that we call *extended multitouch* (EMT) which includes not

only sensing multiple points of touch on and above a surface using a depth camera, but also (1) recovering finger, wrist and hand posture of the user, as well as (2) differentiating users and (3) their handedness. We have implemented a Windows utility that interfaces with a Kinect camera, extracts the above data from the depth image in real-time, and publishes the information as a global hook to other applications, or using the TUIO [16] protocol. To validate our approach, we have also conducted two empirical user studies: one for measuring the tracking accuracy of the touch sensing functionality, and the other for measuring the finger and hand posture accuracy of our utility.

The ability to not only turn any physical surface into a touch-sensitive one, but also to recover full finger and hand posture of any interaction performed on the surface, opens up an entirely new design space of interaction modalities. We have implemented two separate touch applications for the purpose of exploring this design space. The first is a prototype scratchpad where multiple users can draw simple vector graphics using their fingers, each finger on each user being separately identified and drawing a unique color and stroke. The second application is called uSKETCH, and is an advanced sketching application for creating beautified mechanical engineering sketches that can be used for further finite element analysis. uSketch integrates pen-based input that we use for content creation with our touch sensing functionality that we use for navigation and mode changes.

RELATED WORK

Our work is concerned with touch interaction, natural user interfaces, and the integration of several input modalities.

Multitouch Technologies

Several different technologies exist for sensing multiple touches on a surface, including capacitance, RFID, and vision-based approaches. Vision-based approaches are particularly interested for our purposes since they are suitable for surfaces with large form factor, such as tabletops and walls [10, 19, 29]. These systems rely on the presence of a dedicated surface for interaction.

Recently, depth-sensing cameras are beginning to be used in various interactive applications where any surface can be made interactive. Wilson [30] explored the use of depth sensing cameras to sense touch on any surface. Omni-Touch [11] is a wearable depth-sensing and projection system that enables interactive multitouch applications on everyday surfaces. Lightspace [31] is a small room installation that explores interactions between on, above and between surfaces combining multiple depth cameras and projectors.

Sensing Touch Properties

Going beyond sensing touch and deriving *properties* of the touches requires more advanced approaches. We classify existing approaches for this into vision-based techniques, instrumented glove-based tracking, and specialized hardware.

Vision-based Techniques

Malik et al. [19] used two cameras mounted above the surface to distinguish which finger of which hand touched a

surface, but their system requires a black background for reliable recognition. Wang et al. [25] detect the directed finger orientation vector from contact information in real time by considering the dynamics of the finger landing process. Dang et al. [4] mapped fingers to their associated hand by making use of constraints applied to the touch position with finger orientation. Freeman et al. [7] captured the user's hand image on a Microsoft Surface to distinguish hand postures through vision techniques. Holz et al. [14] proposed a generalized perceived input point model in which fingerprints were extracted and recognized to provide accurate touch information and user identification. Finally, Zhang et al. [35] used the orientation of the touching finger for discriminating user touches on a tablet.

Glove-based Techniques

Gloves are widely popular in virtual and augmented reality environments [22]. Wang et al. [26] used a single camera to track a hand wearing an ordinary cloth glove imprinted with a custom pattern. Marquardt et al. [20] used a fiducial-tagged gloves on interactive surfaces to gather input about many parts of a hand and to discriminate between one person's or multiple peoples' hands. Although simple and robust, these approaches require instrumented gear which is not always desirable compared to using bare hands.

Specialized Hardware

Specialized hardware is capable of deriving additional touch properties beyond the mere touch points. Benko et al. [1] sensed muscle activity through forearm electromyography to provide additional information about hand and finger movements away from the surface. Lopes et al. [17] integrated touch with sound to expand the expressiveness of touch interfaces by supporting recognition of different body parts, such as fingers and knuckles. TapSense [12] is a similar acoustic sensing system that allows conventional surfaces to identify the type of object being used for input such as fingertip, pad and nail. The above systems require additional instrumentation of either the user or the surface and take a significant amount of time for setup. Background noise is also a deterrent to use acoustic sensing techniques.

Pen and Touch-based Interaction

A few research efforts have explored the combination of pen and touch [2, 8, 34, 33]. Hinckley et al. [13] described techniques for direct pen + touch input for a prototype, centered on note-taking and scrapbooking of materials. Lopes et al. [18] assessed the suitability of combining pen and multi touch interactions for 3D sketching. Through experimental evidence, they demonstrated that using bimanual interactions simplifies work flows and lowers task times, when compared to the pen-only interface.

OVERVIEW: EXTENDED MULTITOUCH INTERACTION

Traditional multitouch interaction [5] is generally defined as the capability to detect multiple points of touch input on a surface. Here, we enrich this concept into what we call *extended multitouch interaction* (EMT), defined as follows:

- Detecting multiple touch points on and above a surface;

- Recovering finger, wrist and hand postures as well as the handedness (i.e., whether using the left or right hand); and
- Distinguishing between users interacting with the surface.

While previous work has explored various combinations of the above, we are aware of no single work that has addressed all three constraints as a unified framework. To this effect, we propose a two-dimensional hand model that uniquely maps touch points to fingers, hands, and users using depth data. In the following text, we describe our approach to supporting extended multitouch using low-cost depth cameras.

TOUCH SENSING USING A DEPTH CAMERA

Inspired by Wilson [30], our approach to sensing touch using a depth camera simply builds on analyzing the depth image to detect when objects (i.e., fingers) come in close proximity with the physical surface. Our implementation uses a depth camera mounted above a horizontal surface to transform user gestures on and above the tabletop to multitouch events that can be used by a client application. Fully emulating existing multitouch technology using this setup gives rise to several concrete requirements for our system:

- Determine when the surface is being touched and generate the corresponding touch points;
- Track touch points over time when the user is dragging fingers across the surface;
- Map touch points from the physical surface to coordinate space of the display area; and
- Export multitouch events to clients, either using a global Windows hook, or using the TUIO protocol [16].

Below we discuss the physical setup and each of the above components in more detail.

Physical Setup

Our low-cost setup (Figure 2) consists of (1) a horizontal surface, on and above which interactions take place, (2) a projector, which projects on to a surface providing visualization and feedback, (3) a depth sensing camera (like Microsoft Kinect) for capturing touch input, and (4) a computer for processing the input data from the depth camera and projecting visual output using the projector. Again, the surface to be made interactive need not be empty nor flat.

While our focus in this paper is on horizontal (tabletop) setups, there is nothing to prevent the same technique from being used for vertical (wall-mounted) surfaces. The algorithms described below are independent of the configuration and the touch results are not affected as long as a clear line of sight is maintained between the hands and the Kinect.

Basic Touch Sensing

We detect touches by analyzing the image produced by the depth camera. The pixel intensities in the depth image corresponds to the distance from the camera. This allows us to define a “touch input region” as a virtual volume located within a lower and upper distance above the surface. This volume is approximately a cuboid (assuming the surface is

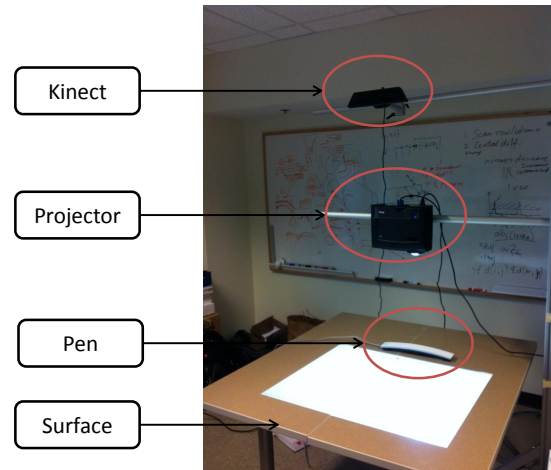


Figure 2. The physical setup showing (from top to bottom) depth camera, projector, pen tracker, and physical surface.




Figure 3. Identifying touch points from the depth data.

flat) offset from the surface with the height equal to the difference of the bounds. All pixels with distance values between these bounds correspond to the contact region(s) of the finger(s) with the surface.

Estimating the underlying surface is important for the purpose of obtaining reliable and robust touches using our method. We model the surface as a depth image that can simply be subtracted from the current image to yield the difference. This difference map is then used to detect whether any physical objects in the scene fall within the touch input region. We create a binary image such that the pixels corresponding to these elements are set to white and all the other pixels to black. Figure 3(a) shows the raw depth image with two hands touching the surface and Figure 3(b) is the binary image showing groups of white pixels corresponding to contacts of fingers with the surface.

We use connected component analysis to identify the cluster of pixels from the touch images, yielding touch “blobs.” To differentiate between actual touches and noise, we discard the blobs that are below or above a certain area threshold. From experimentation, when the Kinect is at a height of 1.2m from the surface, for a blob to be a candidate for a touch point, the area of the blob must be greater than 20 square millimeter and less than 220 square millimeter, accommodating a wide variety of finger sizes. We use the center of gravity of the white blobs as the location of touch points. Figure 3(c) shows the result of determining the touch points from the raw depth data and the binary touch image.

Note that this approach also lends itself to defining additional volumes at other distances from the surface, allowing us to detect interaction above the surface as well.

Tracking Touch Points

Tracking touch points over time is critical for developing interactive applications that make use of dragging gestures. Without this information, each new depth frame would treat touch points as being new, and the application would not be able to support smooth dragging operations.

In our implementation, we associate touch points in consecutive depth frames by finding the nearest neighbor between points in the frames. Since the Kinect can record at a maximum rate of thirty frames per second, not all depth frames are necessarily used in tracking touches. Our system intelligently discards outlier depth frames by considering the touch points in a sliding window that includes several frames before the current frame, thereby increasing robustness.

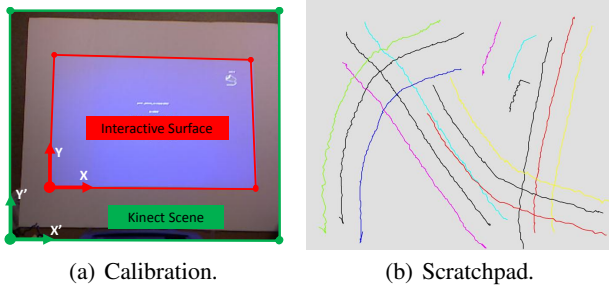


Figure 4. Touch sensing mechanism.

Mapping Touches to Display Space

Traditional multitouch surfaces have a one-to-one mapping between the physical location that is being touched on the surface and the digital location on the underlying display. This provides users with the feeling of directly touching and manipulating the digital content underneath their fingers.

To facilitate a similar interactive experience in our depth camera-based approach, we need to map the depth camera scene to the projected display area (Figure 4(a)). We employ a technique very similar to how regular touch screen sensors are calibrated with LCD displays using a nine-point calibration grid. This gives us the coefficients needed to transform from camera coordinates to display coordinates while taking scaling, rotation, and translation errors into account.

Generating Multitouch Events

The final step in our implementation is to actually generate multitouch events that a client application can handle in order to support (extended) multitouch interaction. We provide two ways to do this: either through a Windows global hook, or using the TUIO protocol [16]. This enables our implementation to be entirely standalone, and external client applications simply need to interface to the system to receive touch events. Figure 4(b) shows a scratchpad application that demonstrates multitouch tracking using these events.

BEYOND STANDARD MULTITOUCH INTERACTION

Our work so far has concerned itself with detecting the two-dimensional (x and y) locations of touch points on as well as above (i.e., with a z component) the physical surface. However, fully supporting our extended multitouch (EMT) model requires being able to recover finger posture, hand posture, and handedness for all touch points, as well as differentiating users. More powerful and rich interaction techniques can be built if such additional information can be leveraged.




Figure 5. Additional touch regions for extended multitouch.

Extracting Palm and Wrist Data

Similar to the touch input region defined earlier for detecting fingers coming into contact with the physical surface, we can extract the position of the user’s palm and wrists from the same depth data. We simply include two additional virtual volumes: the “fingers-palm region” and the “wrist region” (Figure 5). Based on the depth camera image, we construct binary images for each of these regions (Figure 6). These images are then used to recover the touch posture.

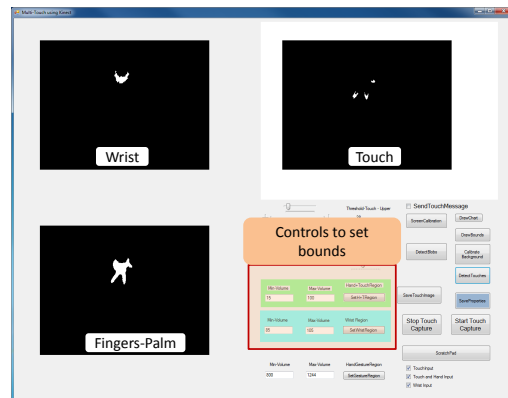


Figure 6. Binary depth images for the three virtual volumes used in our prototype implementation.

A connected component analysis is then performed on the binary image corresponding to ‘fingers-palm region’ to ex-

tract the blobs. For each blob, we compute its a) center of gravity, b) principal directions using principal component analysis and c) convex hull. The center of the palm is approximately the center of gravity of the blob, which is sufficient to map fingers and hands to touch points.

The binary image corresponding to the wrist region is processed in the same manner as above. For each extracted blob, we compute its a) center of gravity and b) principal directions. Wrist orientation can be determined from the principal directions. Tabletop surfaces can take advantage of this information to eliminate the inherent problems that arise due to orientation dependency. Furthermore, this position and orientation information provides an estimate of the hand occlusion region, which can be used to build occlusion-aware interfaces that minimize the impact of occlusion [15, 23, 24]

A Two-Dimensional Hand Model

After processing the three binary images, we combine the results and create a two-dimensional model of the hand. Each hand now has a wrist, fingers-palm, the touch points, and their associations. It can also be noted that the distance value of touches, wrist-center, and palm-center can be obtained from the depth data, which, when added to the two-dimensional model, gives us a three-dimensional model.




Figure 7. From left to right, each image shows the binary image computed in the touch region, fingers-palm region, and the hand model with touch points, wrist and palm.

Figure 7 shows results from two-hand touches. The red dots on the far right correspond to the touch points, the green dot is the palm center, and the violet dot is the wrist. The blue closed polygon is the convex hull of the fingers-palm blob and the green line represents the first principal component of the wrist showing hand orientation. The principal components of the fingers-palm are not shown to avoid clutter. Figure 1 shows 50 touch points from 10 hands and 5 users.

Modeling Touch Posture

To deduce the mapping between the touch points to the fingers and hand, we take advantage of the geometric properties that can be computed from the abstract hand model derived above. In Figure 8, the purple dot is the center position of the wrist and the purple line represents its first principal component. Similarly, the green dot and green line represent the center position and the first principal component of the fingers-palm region. The yellow arrow (V3) represents the direction of the fingers from the wrist passing through the palm. For every instance of the 2D model, we compute (a)

the angles subtended by each touch point with respect to V3, and (b) the distance of each touch point from the wrist.




Figure 8. The angles and distances of the hand geometry.

Wobbrock et al. [32] outlined 27 different gestural commands performed on a tabletop surface. From studying these hand poses, we collected the following observations:

- The majority of gestures are performed with one, two, three, or five finger touch points;
- One-finger interactions are always performed using the fore-finger (100%);
- For two-finger interactions, the fingers used are thumb-fore finger and fore-middle fingers;
- For three-finger interactions, the finger combinations used are thumb-fore-middle and fore-middle-ring; and
- For four-finger interactions, the finger combinations are thumb-fore-middle-ring and fore-middle-ring-and-little.

Based on these observations, we propose 16 different touch postures (Figure 9) to be recognized by our system. Detecting handedness, i.e., if it is a right or left hand, can be determined from the location of the thumb, i.e., its signed angle with respect to the orientation vector. If the signed angle is negative, then the input pattern is a right hand, and vice versa. The classification for recognizing touch posture is based on the number of touch points, which is as follows:

- One - Index finger and ambiguous hand.
- Two - We deduce that a thumb is present if the distance between the two fingers (index and thumb or index and middle) is greater than smallest distance of the fingers from the wrist.



Figure 9. Hand poses during multitouch recognized in our system.




Figure 10. Results (actual snapshots) after mapping touch points to fingers and hands for the poses in Figure 9.

- Three - We deduce that a thumb is present if the angle subtended between extreme fingers (thumb and middle or index and ring) is at least two times greater than the smallest angle made by the extreme fingers with the finger between them (i.e., index and middle or middle and ring finger).
- Four - We deduce that a thumb is present if the angle subtended between the extreme fingers (thumb and ring or index and little) is at least three times greater than the smallest angle made by the extreme fingers with the finger next to them (i.e., ring and little or middle and ring finger).
- Five - the angle and distance between thumb and index fingers is greater than the values between little and ring fingers. Other fingers can now be labeled logically.

Figure 10 shows a sample of the results generated by our system for these 16 different postures. Poses that do not belong to the above classification (for example, Figure 7a) are ambiguous and hence our system cannot identify the fingers reliably. Furthermore, the presence of the thumb and fore finger is crucial in our method for the correct mapping between fingers and touch points, and the thumb is crucial for deducing handedness. Hence, the mapping accuracy decreases for poses which do not have the thumb. In such cases, the distance and angular measures alone do not give conclusive evidence of the correct touch posture and handedness.

Differentiating Users

Finally, being able to match actual touch inputs to individual users is crucial for more advanced collaborative applications where multiple participants are interacting with the tabletop. Our solution builds on the fact that user arms are also partially visible to the depth camera mounted above the tabletop. We use this information to heuristically model user positions that are off-screen of the depth image. However, this approach depends on users maintaining their relative positions around the tabletop surface, as well as not approaching too closely to each other. Also note that our approach does not allow for *identifying users*, only differentiating them.

A future improvement to this solution would integrate a second vision-based tracking system that could follow the participants themselves instead of merely their touch points.

USER STUDIES

To validate our approach, we conducted two empirical user studies: one for measuring the tracking accuracy of the touch sensing functionality, and the other for measuring the finger and hand posture accuracy of our utility.

Study 1: Touch Sensing Accuracy

The accuracy of touch contacts depends on (a) the parameters computed in calibrating the interactive area (projected imagery) to the Kinect scene, (b) the algorithm to detect touch contacts based on depth thresholds, (c) the orientation and size of the finger(s) making the contact, and (d) the user's perception of where they are touching the surface. The above factors all contribute towards the accuracy of the detected touches and as some of them are not controlled, these factors model the real-world performance of our system.

Participants: We recruited 14 volunteers (11 males and 3 females) between the ages of 21 and 32 years for the experiment. Among them, 12 were right-handed and 2 were left-handed.

Task: The task consisted of touching 15 randomly generated locations that were shown on the touch surface. Each location was drawn as a circle with a cross-hair of size 20×20 mm. Each participant was asked to touch every point and hold for two seconds using their index finger. For every location, the system acquired the calibrated touch points repeatedly, between a finger-down and a finger-up event. The system showed one location at a time, and on every finger-up event, a new point was displayed until the end of the experiment.

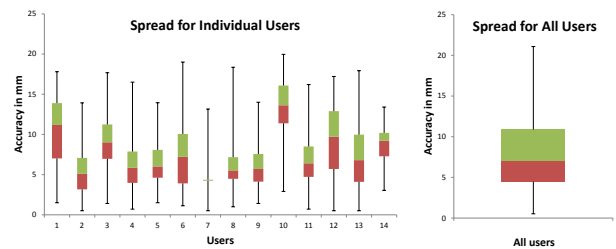


Figure 11. Touch sensing accuracy with average spread for individual users and all users collectively.


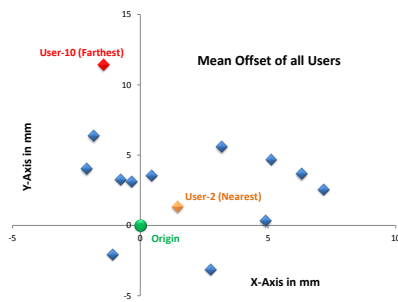
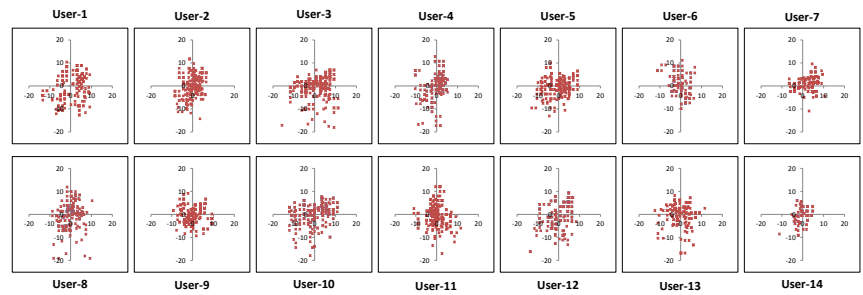


Figure 12. Accuracy of hitting a circular virtual button.

Results: In all, 3715 data points were collected from 14 participants. We analyzed the targeting performance using two



(a) Mean offsets from target (all users).



(b) Distribution of user clicks with compensating mean offsets (mm).

Figure 13. Distribution diagrams.

independent measures: mean offset and spread [11, 14]. Figure 11 shows the average spread for the data collected. The average spread in the best case is 4mm and for the worst case is 13mm. The average spread, taking into account all data points collectively from all users, is 8mm, which is slightly better or in accordance with previous findings in the literature [6, 11]. The mean offset for each user is shown in Figure 13(a) with User-10 having the farthest (largest) offset and User-2 being the nearest (small offset) to the target.

Figure 13(b) shows the distribution of user-clicks for all 14 users with compensating offsets. Here, the origin represents the target locations and the red markers represent the acquired touch points. There were no significant difference in results between the handedness of the users. We also calculated the accuracy of hitting a circular virtual button by accommodating 95% of the touches for each user with compensating offsets. Figure 12 shows the minimum button diameter for every user, with lowest (diameter = 12mm) for User-7 and highest (diameter = 21mm) for User-1. The average diameter for all users is 18mm, which is better than a similar experiment (28mm) conducted in [6]. In comparison, the size of a finger pad is between 10 - 14mm; the fingertip is 8 - 10 mm [3]; the recommended button size for Microsoft Surface is 12mm and for conventional touch input is 15mm [14]. We posit that the accuracy of the touches will decrease with increase in height of the Kinect from the surface. Future studies are needed to study this aspect.

Study 2: Touch Posture and Handedness Accuracy

We conducted a second user study to evaluate the accuracy of recovering touch posture and handedness from depth data.

Participants: We recruited 9 volunteers (7 males and 2 females), between the ages of 22 and 34 for this experiment.

Task: Each participant was asked to perform 8 touch-poses (Figure 9) on the surface with both their left and right hands one after the other. Each pose was to be held for approximately 10 seconds. The participants were not restricted to any position or orientation around the table. No visual feedback was provided to the participant during the study.

Results: A total of 4971 data points were collected during the experiment. The system recognized a hand posture for 4918 of those (approx. 99%). The possible reason for less than 100% recognition is that the users moved their hands

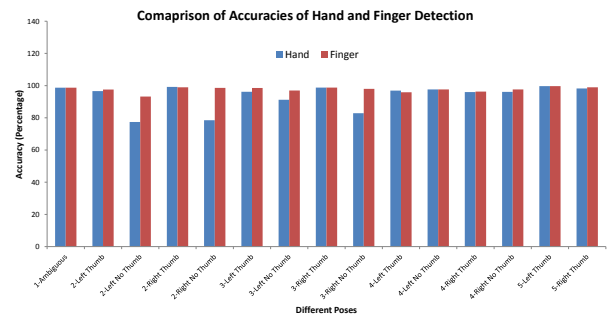


Figure 14. Touch posture and handedness accuracy.

either too fast, or that their hands were not in line of sight to the Kinect depth camera, causing data to be lost.

Figure 14 shows the identification accuracies of touch-posture and handedness for the 16 (8 poses \times 2 hands each) different touch poses. The overall handedness identification accuracy is 93.2%. We discarded one-finger index pose, as our current implementation cannot distinguish the hands for this particular pose. While the two-finger and three-finger “no thumb” poses had an accuracy of 77% and 86% respectively, their corresponding “thumb” poses had an accuracy of 98% and 97% respectively. Thus, the presence of the thumb is critical in identifying the handedness for two-finger and three-finger poses using our algorithm. Our hand model can be improved by reasoning with the contour of the posture and also convexity defects to better identify the posture. However, this will involve more computation and it is important to consider its trade-offs in real-time applications.

The overall finger identification accuracy was 97.7%. At the end of the user study, each participant was asked to rate their least favorable touch-poses. Among all, seven participants voted the four-finger poses (“thumb”, “no-thumb”, left and right hands) and 5 participants voted the “three-finger-no-thumb” pose as their least favorites.

DESIGN IMPLICATIONS

Surfaces supporting extended multitouch (EMT) enable a richer and more diverse interaction design space than those merely supporting standard multitouch. Here we briefly review some of these designs in more detail.

Interaction above and beyond the surface. Existing research has already explored interaction on and above a surface, but we think there is more work to be done here. Touch screens traditionally do not support hovering, but this is certainly one of the straightforward ideas to explore. For example, hovering over a surface could bring up a tool palette, a context menu, or simply a tooltip, similar to how hovering is handled in current mouse-based interfaces. More complex modalities are certainly possible, including detecting volumetric and 3D gestures, as suggested by Wigdor et al. [28].

Identifying users. The ability to match touches to users opens the door for personalized touch interaction. For example, instead of relying on social protocols [21], an EMT application can keep track of object ownership and explicitly prevent users lacking permission from interacting with the object. In a command and control setting, this could be utilized to assign different roles to different participants, such as a fire chief having permission to give orders to fire trucks, whereas a police chief instead may only control squad cars.

However, our extended multitouch implementation currently only supports user differentiation, and identifying users is a more challenging proposition. A heuristic solution may simply assume that particular users reside in different locations around the tabletop surface, allowing them to be identified by their physical position around the table alone.

Recognizing posture. Our implementation has proved the validity of recovering finger and hand posture using a very simple and low-cost hardware setup. This in turn brings several of the posture-based ideas devised by other researchers into the realm of possibility: examples include context-sensitive touch menus, complex gestures on and above the surface, and providing explicit gestural feedback during training. We look forward to seeing what the future will bring in terms of exciting new applications of this new touch modality.

EXAMPLE: PENS WITH EXTENDED MULTITOUCH

Although natural, touch input is not suited for tasks such as writing and sketching. These inking tasks are instead particularly suited for pen-based input, which generally has high precision, whereas our multitouch accuracy is approximately 8mm. Multitouch input in general on all forms and modalities suffer from low precision and our system is no exception. Using a pen for input, on the other hand, is a familiar tool that exploits the users' experience of being trained to use a pen right from their early days.

However, a digital pen realistically allows for only one input point for interaction, constraining the input bandwidth. Augmenting pen-based input with extended multitouch capabilities increases the users' input bandwidth, which can be leveraged by designers to create rich and actually natural interactions for a variety of applications. We outline a division of labor between pen and touch input based on the theme "the pen writes and touch manipulates" [13]. Building on the extended multitouch technique proposed in this paper, we present uSKETCH, an advanced sketching application for creating beautified mechanical engineering sketches that can be used for further finite element analysis.

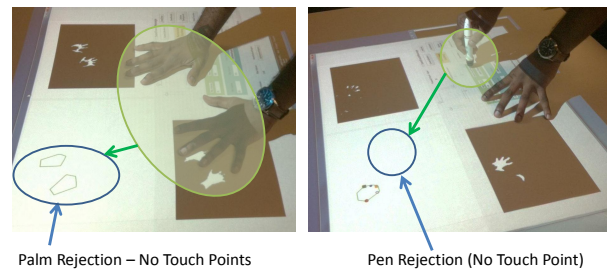


Figure 15. Palm (left) and pen (right) rejection during touch capture.

Combining Pen and Touch Input

For enabling bimanual pen and touch interactions, we added the MimioTM pen to the setup described in Figure 2. Mimio pen technology tracks the 2D position of a digital stylus using ultrasound and infrared, and wirelessly transmits the location to the computer when the stylus is pressed against any flat surface. The pen can be calibrated to any surface size and works both horizontally and vertically. The pen input can then be processed and visually displayed back on the surface using the projector. This process allows the users to write and sketch on any surface, thus making it interactive.

Any object that touches the surface is identified as a touch point by the Kinect sensor. Hence, the system must be able to negate detecting the contact point of the pen as a touch input (Figure 15). To achieve this, we take advantage of how Kinect senses touch. When the pen is in contact with the surface, it communicates its coordinates to the computer. We take advantage of this coordinate information to negate any touch points detected by the system within a small radius of the pen contact. Similarly, when the whole hand rests against the surface, the system detects a big white blob from the palm's footprint on the surface (Figure 15). The system discards the input since it does not correspond to a touch point. Hence, both pen and palm are conveniently discarded by the design of how the system processes the Kinect data.

uSketch

The right part of Figure 1 shows a snapshot of the uSketch application on the interactive surface, and Figure 16 shows the uSketch interface with (a) a freehand sketch of a 2D-bracket, (b) the beautified sketch with interpreted symbols, and (c) the deformation results in ANSYSTM.

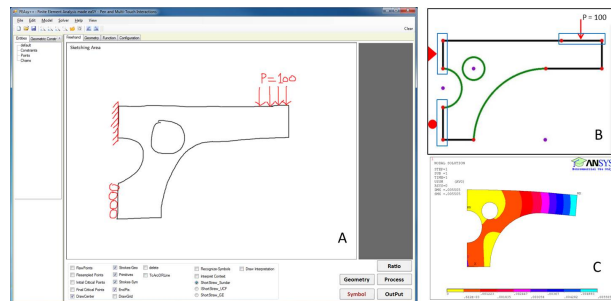


Figure 16. The uSketch interface showing (a) a freehand sketch of a 2D-bracket, (b) the beautified sketch with interpreted symbols, and (c) the deformation results in ANSYSTM.

Guiard [9], proposed general principles for asymmetric bimanual actions comprising of three rules: (a) the dominant hand (DH) moves within the frame of reference defined by the non-dominant hand (NDH); (b) the sequence of motion generally sees the NDH setting the reference frame prior to actions with the DH being made within that context; and (c) that the DH works at a higher level of precision than the NDH in both spatial and temporal terms. Hence, the pen is preferably held in the DH and both NDH and DH be used for touch-input with the pen tucked in the DH.

Using bimanual interactions with pen and touch eliminates the need for constant switching between modes and also the need for user interface elements for changing visualizations and other interactions. We posit that using such interactions simplifies work flows and lowers task times, when compared to the pen-only or touch-only interface.

uSketch and EMT

The uSketch tool was designed specifically for taking advantage of extended multitouch. Below we outline some of the supported actions and their motivation:

- **Sketching, symbols, and text:** Pen input only; pens are optimal for content creation [13].
- **Navigation:** One finger from each hand, or one finger and pen; hand is quick and precise enough for navigation.
- **Deletion:** Three fingers on NDH and squiggle with pen or finger on DH; complex gesture to avoid accidental trigger.
- **Context menu:** Three fingers tap or down on NDH, selection using DH pen or finger; complex gesture as above.

Deploying uSketch

We designed uSketch as a design-analysis tool that introduces new ways for mechanical engineering students to interact, explore, and learn design and analysis in the classroom and in the laboratory. The tool provides a complete workflow from early sketching, through beautified imagery, and finally to finite element analysis, and is designed to teach a wide variety of mechanical engineering concepts in courses such as statics, strength of materials, and kinematics. We are currently deploying the tool in a toy design course (Purdue ME 444) where we use it to teach students the connections between early design (creativity) and computer-aided design (final product).

By supporting for hypotheses, calculations, and discussions, uSketch builds a strong intuition for engineering design decisions among students beyond the static content in textbooks.

CONCLUSIONS AND FUTURE WORK

We have presented a multitouch sensing approach based on low-cost depth cameras where any physical surface, flat or non-flat, horizontal or vertical, projected or screen, can be instrumented with touch sensing. Our contribution goes beyond existing depth camera touch sensing approaches into what we call *extended multitouch*, which we define as the ability to not only sense multiple touches, but also detecting interaction above the touch surface, recovering the user's finger and hand posture, and distinguishing between the users interacting with the surface. We validate the work using

two user studies gauging the accuracy of touch and posture recognition, as well as using two sketch applications where we begin to explore the design space of this novel idea.

We have only scratched the surface of this extended multi-touch interaction approach, and, as discussed in the design implications above, we can see many opportunities for future work in this area. In particular, we envision continuing our investigations into combining multiple interaction modalities for early engineering design and sketching.

ACKNOWLEDGMENTS

This work is partly supported by the National Science Foundation Partnership for Innovation Grant No. 0917959, the Donald W. Feddersen Chaired Professorship, and the Purdue School of Mechanical Engineering. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

REFERENCES

1. H. Benko and A. Wilson. DepthTouch: Using depth-sensing camera to enable freehand interactions on and above the interactive surface. In *Proceedings of the IEEE Workshop on Tabletops and Interactive Surfaces*, volume 8, 2009.
2. P. Brandl, C. Forlines, D. Wigdor, M. Haller, and C. Shen. Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*, 154–161, 2008.
3. K. Dandekar, B. I. Raju, and M. A. Srinivasan. 3-D finite-element models of human and monkey fingertips to investigate the mechanics of tactile sense. *Journal of Biomechanical Engineering*, 125(5):682–691, 2003.
4. C. T. Dang, M. Straub, and E. André. Hand distinction for multi-touch tabletop interaction. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, 101–108, 2009.
5. P. Dietz and D. Leigh. DiamondTouch: a multi-user touch technology. In *Proceedings of the ACM Symposium on User interface Software and Technology*, 219–226, 2001.
6. A. Dippon and G. Klinker. KinectTouch: accuracy test for a very low-cost 2.5D multitouch tracking system. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, 49–52, 2011.
7. D. Freeman, H. Benko, M. R. Morris, and D. Wigdor. ShadowGuides: visualizations for in-situ learning of multi-touch and whole-hand gestures. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, 165–172, 2009.
8. M. Frisch, J. Heydekorn, and R. Dachselt. Investigating multi-touch and pen gestures for diagram editing on interactive surfaces. In *Proceedings of ACM Conference on Interactive Tabletops and Surfaces*, 149–156, 2009.
9. Y. Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *J Mot Behav*, 19(4):486–517, 1987.

10. J. Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 115–118, 2005.
11. C. Harrison, H. Benko, and A. D. Wilson. OmniTouch: wearable multitouch interaction everywhere. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 441–450, 2011.
12. C. Harrison, J. Schwarz, and S. E. Hudson. TapSense: enhancing finger interaction on touch surfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 627–636, 2011.
13. K. Hinckley, K. Yatani, M. Pahud, N. Coddington, J. Rodenhouse, A. Wilson, H. Benko, and B. Buxton. Manual deskterity: an exploration of simultaneous pen + touch direct input. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, 2793–2802, 2010.
14. C. Holz and P. Baudisch. The generalized perceived input point model and how to double touch accuracy by extracting fingerprints. In *Proceedings of ACM Conference on Human Factors in Computing Systems*, 581–590, 2010.
15. W. Javed, K. Kim, S. Ghani, and N. Elmqvist. Evaluating physical/virtual occlusion management techniques for horizontal displays. In *Human-Computer Interaction - Proceedings of INTERACT*, volume 6948 of *Lecture Notes in Computer Science*, 391–408. Springer, 2011.
16. M. Kaltenbrunner. reacTIVision and TUIO: A tangible tabletop toolkit. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, 9–16, 2009.
17. P. Lopes, R. Jota, and J. A. Jorge. Augmenting touch interaction through acoustic sensing. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, 53–56, 2011.
18. P. Lopes, D. Mendes, B. Araújo, and J. A. Jorge. Combining bimanual manipulation and pen-based input for 3D modelling. In *Proceedings of the Eurographics Symposium on Sketch-Based Interfaces and Modeling*, 15–22, 2011.
19. S. Malik and J. Laszlo. Visual touchpad: a two-handed gestural input device. In *Proceedings of the ACM International Conference on Multimodal Interfaces*, 289–296, 2004.
20. N. Marquardt, J. Kiemer, D. Ledo, S. Boring, and S. Greenberg. Designing user-, hand-, and handpart-aware tabletop interactions with the TouchID toolkit. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, 21–30, 2011.
21. M. R. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond 'social protocols': multi-user coordination policies for co-located groupware. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 262–265, 2004.
22. D. J. Sturman and D. Zeltzer. A survey of glove-based input. *IEEE Computer Graphics and Applications*, 14(1):30–39, Jan. 1994.
23. D. Vogel and R. Balakrishnan. Occlusion-aware interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 263–272, 2010.
24. D. Vogel and G. Casiez. Hand occlusion on a multi-touch tabletop. In *Proceedings of ACM Conference on Human Factors in Computing Systems*, 2307–2316, 2012.
25. F. Wang, X. Cao, X. Ren, and P. Iran. Detecting and leveraging finger orientation for interaction with direct-touch surfaces. In *Proceedings of ACM Symposium on User Interface Software and Technology*, 23–32, 2009.
26. R. Y. Wang and J. Popović. Real-time hand-tracking with a color glove. *ACM Transactions on Graphics*, 28(3), 2009.
27. M. Weiser. The computer for the twenty-first century. *Scientific American*, 3(265):94–104, Sept. 1991.
28. D. Wigdor, H. Benko, J. Pella, J. Lombardo, and S. Williams. Rock & rails: extending multi-touch interactions with shape gestures to enable precise spatial manipulations. In *Proceedings of ACM Conference on Human Factors in Computing Systems*, 1581–1590, 2011.
29. A. D. Wilson. TouchLight: an imaging touch screen and display for gesture-based interaction. In *Proceedings of the ACM International Conference on Multimodal interfaces*, 69–76, 2004.
30. A. D. Wilson. Using a depth camera as a touch sensor. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, 69–72, 2010.
31. A. D. Wilson and H. Benko. Combining multiple depth cameras and projectors for interactions on, above and between surfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 273–282, 2010.
32. J. O. Wobbrock, M. R. Morris, and A. D. Wilson. User-defined gestures for surface computing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 1083–1092, 2009.
33. K.-P. Yee. Two-handed interaction on a tablet display. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*, 1493–1496, 2004.
34. R. Zeleznik, A. Bragdon, F. Adeptura, and H.-S. Ko. Hands-on math: a page-based multi-touch and pen desktop for technical work and problem solving. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 17–26, 2010.
35. H. Zhang, X.-D. Yang, B. Ens, H.-N. Liang, P. Boulanger, and P. Irani. See me, see you: a lightweight method for discriminating user touches on tabletop displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2327–2336, 2012.