Typing on Flat Glass: Examining Ten-Finger Expert Typing Patterns on Touch Surfaces

Leah Findlater,[1] **Jacob O. Wobbrock,**[1] **Daniel Wigdor**[1,2]

[1] The Information School
University of Washington
{ leahkf, wobbrock }@uw.edu

[2] Microsoft Research
Redmond, WA
dwigdor@microsoft.com

ABSTRACT

Touch screen surfaces large enough for ten-finger input have become increasingly popular, yet typing on touch screens pales in comparison to physical keyboards. We examine typing patterns that emerge when expert users of physical keyboards touch-type on a flat surface. Our aim is to inform future designs of touch screen keyboards, with the ultimate goal of supporting touch-typing with limited tactile feedback. To study the issues inherent to flat-glass typing, we asked 20 expert typists to enter text under three conditions: (1) with no visual keyboard and no feedback on input errors, then (2) with and (3) without a visual keyboard, but with some feedback. We analyzed touch contact points and hand contours, looking at attributes such as natural finger positioning, the spread of hits among individual keys, and the pattern of non-finger touches. We also show that expert typists exhibit spatially consistent key press distributions within an individual, which provides evidence that eyes-free touch-typing may be possible on touch surfaces and points to the role of personalization in such a solution. We conclude with implications for design.

Author Keywords

Touch-typing, multi-touch input, tabletop computing.

ACM Classification Keywords

H.5.2. [Information interfaces and presentation]: User interfaces—input devices and strategies.

General Terms

Human factors, design, experimentation.

INTRODUCTION

Touch surfaces have become pervasive, most commonly in the form of mobile phones, and increasingly as larger tablets and interactive tabletops (e.g., Apple iPad and Microsoft Surface). Unlike their mobile phone counterparts, tablets and tabletops are large enough to accommodate ten-finger text input. Despite the importance of text entry on these devices, however, common methods for entering text pale in effectiveness compared to typing on a physical keyboard [2,12]. Indeed, text entry has been called out as a major deficiency of tabletops in multiple studies (e.g. [3,25]), potentially detracting from sustained use [22,32] and even leading users to alter their linguistic style [32]. Physical keyboards have been proposed as a solution for tabletop interaction [10], which, although useful, highlights the failings of current software-based methods.

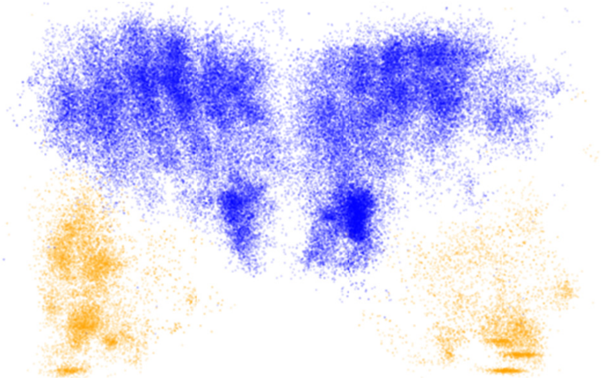

Figure 1. Finger (blue) and non-finger (yellow) touches in the unrestricted typing condition, showing space between hands, separate left and right space bar areas, and evidence of forearms and heels of the hands resting on the screen. (*N* = 20)

A major challenge of typing on flat surfaces is the lack of tactile feedback, not just in the physical perception of key depressions but also in the ability to feel the home row [2]. As a result, users must visually attend to the keyboard at all times. Inadvertent touches occur frequently on tabletops [25], which makes touch screen keyboards vulnerable to spurious input, especially from the user's arm or palm. However, touch screen keyboards also offer rich potential for customization and adaptation because they are software-based. Despite this potential, commercial implementations of ten-finger keyboards have largely mimicked the traditional QWERTY keyboard (e.g., Apple iPad, Microsoft Surface, Canesta keyboard [24]). Research approaches include adaptive keyboards [5,11] and wearable keyboards [6], but have not yet offered high performance text input.

In this paper, we examine typing patterns that emerge when expert users of physical keyboards touch-type on a flat surface (Figure 1). Our aim is to inform future designs of touch screen keyboards, with the ultimate goal of supporting touch-typing with limited tactile feedback. We employed a technique evocative of a *user-defined interface* (e.g., [33]): we asked 20 touch-typists to enter phrases using

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

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
Copyright 2011 ACM 978-1-4503-0267-8/11/05....$10.00.

an interactive tabletop under three conditions: (1) with no visual keyboard and no feedback, then (2) with and (3) without a visual keyboard, and with some feedback. We were interested in questions such as: What will the distribution of key presses look like if users are given no visual constraints when typing? Will the centroids for each key follow the layout of a standard rectangular keyboard? Will certain keys have a larger spread of hits? By using vision algorithms to detect hand contours, can we identify differences among users in key-to-hand mappings?

Our findings have implications for the design of static and personalized touch screen keyboards. Based on finger touch data, we show that a curved keyboard with a gap between the hands is a more natural representation of actual typing patterns on flat surfaces than a standard rectangular keyboard. We also show that some keys are more difficult to hit consistently than others, suggesting those keys should be made larger (e.g., keys assigned to the little fingers). Typing patterns varied widely among users, but finger placement per key was highly reliable within an individual: with a simple classification approach using centroids of the key hit points, we classified key presses at 90% accuracy in a condition where there was *no visual keyboard*.

This paper contributes a formative study of unconstrained typing patterns on a flat surface, and an empirical basis for future development of ten-finger flat-surface keyboards. We also show that expert typists exhibit spatially consistent key press distributions within an individual, which provides evidence that eyes-free text input may be possible on touch surfaces and points to the role of personalization in such a solution. Finally, we also contribute design implications for both static and personalized touch screen keyboards.

RELATED WORK

Adding to the techniques mentioned in the Introduction, we discuss text entry for large multi-touch devices, virtual keyboard work in general, and physical keyboards.

Tabletops & Text Entry

In 2007, Hinrichs et al. [12] surveyed text entry techniques for tabletops, dividing the space into external methods (e.g., physical keyboards) and on-screen methods (e.g., virtual keyboards). Most of the on-screen methods were for small, mobile touch screens and none supported ten fingers. Hinrichs et al. [12] also offered evaluation criteria for tabletop text entry, such as the need for rotatability, on-screen mobility, and support for multi-person interaction. In an observational study of tabletop use, Ryall et al. [25] offered another general finding that has implications for typing: the difficulty of distinguishing between intentional and inadvertent touches of fingers, hands, and arms.

Complementing our work, recent research has focused on tactile feedback to support tabletop text entry. Weiss et al. [31] proposed a silicon keyboard overlay that could potentially support eyes-free typing; no evaluation has yet been reported. McAdam and Brewster [17] studied distal tactile feedback during text entry and found that feedback on the wrist or upper arm improved typing speed. None of the previous work has studied ten-finger typing patterns, at most reporting speed, error rates, and subjective feedback.

Virtual Keyboards

Virtual keyboards provide a temporary allocation of screen space for text entry. Solutions using a mouse, eye-trackers, or other assistive technologies have been studied for accessible text entry (e.g. [16,28]). For the broader community, studies have largely focused on stylus or direct-touch interfaces. Although findings described in this section are highly relevant to ten-finger typing, there will be differences in biomechanics and efficiency because most of these techniques support a single point of input.

Past research has examined key positioning and size. Sears et al. [27] found that smaller keys reduced text entry speed and increased errors. However, later work by MacKenzie and Zhang [21] found that, although a smaller keyboard increased errors compared to a larger one, there was no reduction in speed. This is in keeping with the application of Fitts' law to performance optimization of virtual keyboards [38].

Researchers have also explored relaxing the requirement to precisely hit each key. Kristensson and Zhai [14] proposed a method whereby the overall geometric shape formed by all of the hit points for a word is considered in linguistic matching. This approach was expanded on by Rashid and Smith [23] to enable typing without a priori determining the position of the keyboard, albeit with an extremely high error rate. Gunawardana et al. [9] proposed a method to expand or contract key areas for each press based on linguistic models, building on previous work by Goodman et al. [7] and Al Faraj et al. [1]. In similar work, Himberg et al. proposed adaptation through the movement of individual keys [11]. Alternatives to tapping a virtual QWERTY keyboard have also been proposed, including alternate key layouts [15,20,35,37], gestures [30], and methods that enable users to stroke between keys [13,36].

Physical Keyboards

The first practical typewriter was introduced in 1874 and touch-typing gained prominence a few decades later due to performance advantages and reduced fatigue over hunt-and-peck typing [4]. The development of touch-typing expertise requires extensive training [4], with skilled typists reaching speeds of 60 WPM or higher [8].

Previous work on common typing errors serves as the basis for some of our own analysis. Errors include misstrokes, which result from inaccurate finger movement, omissions, insertions, and interchanging of letters in the text [8]. Substitution errors, where one letter is substituted for another, occur most commonly in the same row or column, and can even be homologous (mirror-image position on the opposite hand) [8]. In a survey of typing studies, one relevant finding Salthouse [26] reports is that different fingers result in different error frequencies.

STUDY OF SKILLED TYPISTS

To observe unconstrained typing patterns on a flat surface, we asked skilled typists to enter text under three conditions:

1) no feedback and no keyboard (*unrestricted* typing),
2) asterisk feedback and no keyboard,
3) asterisk feedback and a visible keyboard.

The two conditions without a visible keyboard (1-2) captured natural typing patterns without the user adapting to a particular virtual form factor. In the *unrestricted* condition (1), participants were unaware of spurious and missing touches, which allowed for the most natural typing possible, mimicking an ideal touch-typing keyboard. In the *feedback* conditions (2-3), output was in the form of asterisks for each non-space touch. The asterisks allowed users to avoid spurious touches and ensure touches registered with the device without causing changes in pose and finger placement by showing which key they had hit. Feedback was also necessary to create a 1:1 mapping between touch events and letters in the presented phrase in the *asterisk feedback, no keyboard* condition (2), since there was no underlying model of key locations.

Participants

Twenty skilled touch-typists (8 female) ranging in age from 20 to 44 (*M* = 29.3) participated in our study. All participants regularly used a rectangular physical keyboard (not a split or "natural" keyboard). We evaluated typing expertise with a physical keyboard typing test at the start of each session using *TextTest*, a text entry evaluation tool [34]. Participants typed twenty phrases, randomly selected from the MacKenzie phrase set [19] (the same phrase set as used for the experimental tasks). Mean typing speed was 85.0 WPM (*SD* = 19.4) with a mean corrected error rate of 2.5% (*SD* = 1.5) and a mean uncorrected error rate of 0.1% (*SD* = 0.2). Only two participants had experience with an interactive table, but all participants had used a touch screen mobile phone. Participants were compensated $30.

Apparatus

Participants sat at a Microsoft Surface table, which ran the experiment software. The Surface API reported touch events, which served as the input signal. In addition, our software processed raw images provided by the Surface to extract additional information about the participant's hand placement. We computed the convex hull of each hand using custom vision algorithms and Emgu CV, an open source C♯ wrapper for the OpenCV library.[1]

Figure 2 shows a screenshot of the experiment software, with the presented phrase at the top of the screen, a rectangular text input area, and the "Next phrase" button, which allowed participants to advance through the task. The application was coded in C♯ .NET 3.5. The system recorded logging data with millisecond timestamps. For the physical typing test, an external 22" monitor (1680×1050 resolution) and rectangular keyboard were used.

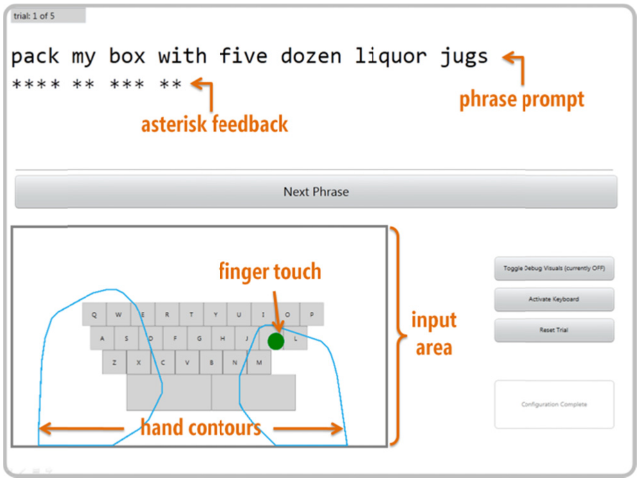

trial: 1 of 5

pack my box with five dozen liquor jugs

**** ** *** **

Figure 2. Task interface, showing the *asterisk feedback, visible keyboard* condition. Hand contours and finger touch point are for illustration only and were not displayed to users.

Figure 3. Input area at start of a trial in the *unrestricted* and *asterisk feedback, no keyboard* conditions. Red dots indicate where to place thumbs based on per-user configuration.

Typing Interfaces

A short *configuration step* was required at the start of each condition: participants placed their hands in a comfortable position within the input area (Figure 2) and the system recorded the (*x*,*y*) locations of the thumbs, and used them to place the keyboard. Participants could place their hands anywhere within the input area for this configuration step. The input area was inactive at the start of each trial.

Condition 1: Unrestricted (No Feedback, No Keyboard)

The goal of this condition was to observe how expert typists position their hands and fingers on the screen and type without the constraints of a visual keyboard layout or issues with touch-input recognition errors. We were interested in studying hand and finger placement, such as whether participants would rest their fingers on the home row during trials. Participants used the interface shown in Figure 2, except that the input area was *completely* blank and *no feedback* was provided on typing input (i.e., no asterisks). Participants were instructed to, "type comfortably and naturally, but also quickly and accurately, as you would normally do on a regular keyboard."

To maintain a degree of consistency across trials, two red dots appeared at the start of each trial, where the participant's thumbs had been during the configuration step (Figure 3). Participants placed their hands in the input area with their thumbs over the red dots and requested that the

---

[1] http://www.emgu.com and http://opencv.willowgarage.com

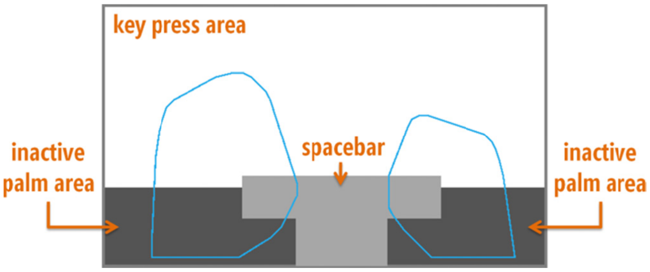

Figure 4. The *asterisk feedback, no keyboard* condition had inactive areas and a T-shaped spacebar that were not visible.

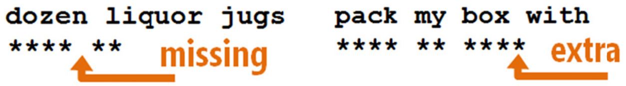

Figure 5. Sample *asterisk feedback*, showing input errors.

experimenter activate the keyboard. The experimenter did so by pressing a button on the right of the screen, causing the red dots to disappear and the input area to turn light blue. The participant entered the phrase by touch-typing on the flat-glass surface. No backspace was provided, but participants could request to restart the trial if they felt they had committed typing errors.

Conditions 2 & 3: Asterisk Feedback

The goal of the *asterisk feedback* conditions was to collect an input stream that could be reliably processed and mapped to user intention. Accurately labeled touch events were necessary to evaluate typing patterns in detail.

Condition 2: Asterisk feedback, no keyboard. Similar to the *unrestricted* condition, this condition presented only a blank input area to participants. Again, based on the location of the thumbs during the configuration step, two red dots appeared at the start of each trial to guide participants to consistent hand placement across trials (Figure 3); the dots disappeared after input area activation. Although only a blank input area was shown to users, unlike the *unrestricted* condition, an underlying keyboard model disambiguated spacebar presses, letter key presses, and all other touches. This allowed spacebar presses to be properly labeled as spaces and key presses as asterisks for the visual feedback. Figure 4 displays elements of the underlying input model that were *not* visible during the test: (1) a T-shaped spacebar area was placed based on the location of thumbs during the configuration step, such that the intersection of the T was at the midpoint between the thumbs, and (2) inactive areas on either side of the spacebar, to reduce spurious key presses from palms or forearms. The simple heuristic used to place the spacebar was enough that the hidden spacebar worked for most participants without having to reconfigure the location of the thumbs.

Condition 3: Asterisk feedback, visible keyboard. This condition presented a rectangular QWERTY keyboard (Figure 2). Key size and placement were based on the touch screen keyboard provided by the Microsoft Surface. Keys were 0.9" in width and height, slightly larger than a physical keyboard (0.75"). We confirmed the use of the larger size based on feedback during early pilots that 0.75" was too small. The space key was taller (by 1.5 times) than the other keys, as is common on physical keyboards. The keyboard was placed such that the center of the spacebar was positioned at the mid-point of the thumb location during initial configuration. This positioning was meant to place the keyboard most comfortably for each user.

Instructions in both *asterisk feedback* conditions were to type comfortably and naturally, but with an emphasis to place fingers accurately and to correct errors. As such, a backspace gesture was provided as a right-to-left swipe (> 0.5") anywhere on the right half of the keyboard, we used the gesture instead of a backspace key to avoid hand drift while correcting errors. Participants were instructed to ensure their asterisks and spaces lined up with characters in the presented text before advancing to the next trial (Figure 5). The ambiguity of the asterisks often made the precise location (i.e., character index) of an error difficult to identify, especially if participants typed quickly. Participants were asked to delete the *entire* word containing an extra or missing character if they were uncertain about the exact location of the error. This obviously affected text entry speeds but gave us the ability to infer users' intention, which is critical for the purposes of our study.

Rather than having the experimenter activate the keyboard as was done in the *unrestricted* condition, participants activated the keyboard themselves by briefly placing 10 fingers on the screen then removing them. This action allowed participants to situate their fingers on home row (whether or not a visible home row existed), and was meant to reduce inadvertent touches at the start of a trial. To provide feedback on the activation, the input area turned orange when all 10 fingers were detected, and blue once the fingers were lifted. Because the goals of the *asterisk feedback* conditions were not the same as for the *unrestricted* condition, the inconsistency in activation method did not affect the quality of collected data.

Experiment Design

We used a within-subjects, single-factor design (*typing interface*). The *unrestricted* condition was always presented first so that participants would not be biased by experience with input errors on the touch screen or by the layout of the visual keyboard. The *asterisk feedback* conditions were counterbalanced, with participants randomly assigned to orders. As such, we treat data from the *unrestricted* and *asterisk feedback* conditions separately.

Procedure

Study sessions were designed to last 1.5 hours. Participants completed a questionnaire to collect demographic information and experience with touch screens, followed by the physical keyboard typing test. Participants then completed 5 practice phrases and 40 test phrases with each of the three typing conditions. Opportunities to rest were provided midway through the test blocks. To cover all letters in the alphabet, half of the phrases in the test set consisted of two pangrams: *the quick brown fox jumped over the lazy dog* and *pack my box with five dozen liquor jugs*, while the remaining phrases were randomly selected from the MacKenzie phrase set [19]. The random phrases and pangrams were intermixed. For each condition, the thumb configuration step was done once for the practice phrases and redone before the test phrases to allow participants to adjust their hands if they wished. Finally, feedback questionnaires were administered.

Exploratory Questions

The three conditions were designed to answer different questions. We summarize the most salient here.

Unrestricted Typing

1. How fast can users type on a flat surface when they assume their input is accurate?
2. How many finger and non-finger touches occur and where do they occur?
3. How do the number of finger touches compare to the length of the presented text?

Asterisk Feedback, Visible Keyboard and No Keyboard

1. How fast will users type with a visible keyboard versus no keyboard? How many errors will they commit?
2. Is the emergent keyboard layout based on actual key presses different between the two conditions in terms of curvature and distance between hands?
3. Do some keys have greater x- or y-axis deviation than others? Are such findings systematic by row or column?
4. Do key-to-hand mappings follow the touch-typing standard (`T`, `G`, `B` to left hand and `Y`, `H`, `N` to right hand)?
5. Are key press locations for each key consistent? How reliably can we classify key presses based on the observed centroids of key presses?

Data and Analysis

Across all participants, we collected 50,289 labeled key presses from the *asterisk* tasks, and 27,830 unlabeled finger touches in the *unrestricted* condition. In addition, on every touch *down* event (as opposed to *moved* or *up* events), we processed the raw image and recorded the convex hull around each hand. Due to a technical problem, only the left hand convex hull was recorded for the first 8 participants—this was remedied for the remaining participants.

Although we asked participants to correct all errors in the *asterisk* conditions, the ambiguity of providing asterisks as output meant that uncertainty remained in the labeling. Expert typists recognize between 40-70% of their own typing errors by feel and without visual feedback [26], which means that errors likely remained in the data due to the ambiguous visual feedback we provided. Clear cases of mislabeled key presses can be identified: for example, typing `E` then `M` instead of the opposite. To account for these mislabelings, we removed outlying points for each key that were more than three standard deviations away from the mean in either the x or y direction (1.8% of instances).

We used repeated measures ANOVAs and paired two-tailed t-tests for our analyses. All *post hoc* pairwise comparisons following the ANOVAs were protected against Type I error using a Bonferroni adjustment. Reported fractional degrees of freedom (*df*s) are from Greenhouse-Geisser adjustments. When parametric tests were not appropriate because the data violated the assumption of normality, we applied non-parametric equivalents, such as the Wilcoxon signed-rank test. We report significant findings at $p < .05$.

RESULTS

We examine the *unrestricted* typing condition before exploring the *asterisk feedback* conditions in more depth.

Unrestricted Typing: No Keyboard, No Feedback

Our goals were to learn how quickly users can type on a flat surface when they assume their input is accurate, and to observe the pattern of touches, especially those that were not the result of user-intended actions. The mean number of reset trials per user was 20.8% ($SD = 9.9$). Figure 1 shows finger and non-finger touch points for all participants.

Typing speed was 31% slower than the physical keyboard. We calculated WPM following MacKenzie [18]:

$$\text{WPM} = \frac{|T| - 1}{S} \times 60 \times \frac{1}{5}$$

where T is the final transcribed string and S is the elapsed time in seconds, in our case, from first to last finger touch in a trial. For reset trials, we discarded input from before the reset. This measure provides an indication of the speed that users could achieve on a flat surface under ideal conditions. Participants typed an average of 58.5 WPM ($SD = 18.0$), with a range of 31.3 to 92.7 WPM. Although mean WPM was 31% slower than the physical keyboard, it is still almost 25 WPM faster than the predicted expert typing speed of a stylus-based QWERTY keyboard [37].

Fingers often rested on the screen, especially at the start of a trial. At the start of each trial, participants placed their hands in preparation for typing, with most participants resting at least some fingers on the screen, as if on the home row (fingers down at start: $M = 5.08$, $SD = 2.94$, range: 0 to 10). We also compared the number of finger touches *after* the trial start to the length of the presented text. There were slightly more finger touches than expected (ratio of finger touches to length of the presented text: $M = 1.07$, $SD = 0.07$), indicating that participants sometimes rested their fingers on the screen or brushed them inadvertently against it. However, looking at occurrences where three or more fingers were down simultaneously, we found that only five participants exhibited this behavior more than once (for those five: $M = 0.45$ occurrences per trial, $SD = 0.48$). These results demonstrate the need to support fingers resting on the home row between text entry sequences, but not necessarily during a sequence.

Hand and arm position varied. Fewer than half the participants ($N = 7$) consistently rested their hands on the screen, while the remaining participants rested their forearms on the edge of the table and hovered their palms. These behaviors resulted in more than one non-finger touch per word (per trial: $M = 8.24$, $SD = 4.41$). As Figure 1 shows, however, finger and non-finger touches were highly segmented, suggesting that a simple approach of creating an inactive area under each palm should be highly effective at filtering non-finger touches (see Discussion for details).

Asterisk Feedback Conditions

From the *unrestricted* typing condition, we learned that users can type quickly on a flat surface if they assume their feedback is correct. With the *asterisk feedback* conditions we can conduct a more detailed analysis of typing patterns using the 1:1 correspondence between touch events and presented text. The *no keyboard* condition provides data to most directly support touch-typing (i.e., eyes-free) on a flat surface. We compare the *no keyboard* and *visible keyboard* conditions in terms of emergent keyboard shape, the spread of hits for individual keys, key-to-hand mappings, and how accurately finger touches can be classified as specific keys. Figure 6 shows key presses for all participants.

Mean trial time was 17.8 seconds (*SD* = 5.0) with the *visible keyboard* and 18.7 seconds (*SD* = 5.4) with *no keyboard*. Participants reset on average 8.1% of trials (*SD* = 9.7) with *no keyboard* and 4.7% of trials (*SD* = 5.0) with a *visible keyboard* (a Wilxocon signed-rank test was not significant: $z = 1.45$, $p = .147$). Analyses in this section are done on the text entered after a reset, if one occurred.

On average, 81.4% of key presses fell within the bounding box of their corresponding key for the *visible keyboard* condition (*SD* = 12.4). This number may seem low, but it is not surprising given that participants were provided only with asterisks and spaces as feedback. We revisit the reliability of key presses later, in the Classification section.

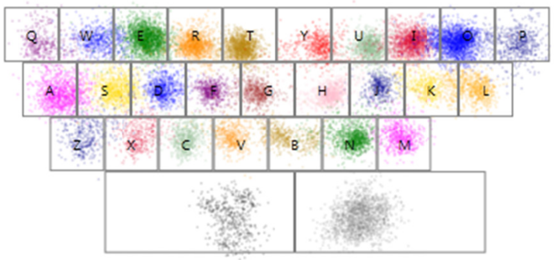

(a) ***visible keyboard***

(b) ***no keyboard*** **(overlay is for illustration only)**

Figure 6. All key presses in *asterisk feedback* conditions, colored by key label. Each participant's data is translated to the same midpoint between F and J. The *visible keyboard* shows more consistency across users than *no keyboard*. (*N* = 20)

Typing Speed and Errors (Keystrokes Per Character)

Although the main goal of the *asterisk feedback* conditions was not to measure typing speed and errors, we analyze them for completeness. We calculated WPM as for the *unrestricted* condition, with elapsed time as the first to last key press per trial. The average WPM was similar across conditions: 27.5 (*SD* = 7.8) with the *visible keyboard* and 28.1 (*SD* = 9.6) with *no keyboard*. This difference was not statistically significant ($t_{19} = 0.43$, $p = .669$). That participants were slower in these trials than in the *unrestricted* typing condition is unsurprising, given the emphasis on accuracy and the presence of feedback.

To quantify errors, we calculated the keystrokes entered per character (KSPC) ratio [29]:

$$\text{KSPC} = \frac{|IS|}{|T|}$$

where *IS* is the input stream (including backspaces) and *T* is the transcribed string. Recall that participants had been asked to delete the entire word containing an error if there was uncertainty about where the error had occurred. As a result, KSPC was relatively high for both conditions: on average, 1.26 (*SD* = 0.02) for the *visible keyboard* and 1.29 (*SD* = 0.03) for *no keyboard*. Since we counted backspace gestures as key presses for this analysis, 13.0% and 14.5% of typed keys were corrected in the *visible keyboard* and *no keyboard* conditions, respectively. Overall, the similarity of the two conditions indicates participants were not significantly hindered by having no visual keyboard to reference.

Overall Keyboard Shape

We were interested in the overall keyboard shapes that would emerge when users were given a visual reference versus having no visual constraints (similar to touch-typing). For this analysis, we calculated key centroids for each participant. Since participants could place their hands anywhere in the input area to type, we normalized the centroids so that each participant's centroids were centered around the midpoint between their observed F and J key centroids (F and J often provide small raised bumps on a physical keyboard, and thus are used to place the hands). Centroids across all participants are shown in Figure 7.

Emergent keyboard shape is more arched in the 'no keyboard' condition than the 'visible keyboard' condition. We examined whether having no visual constraints would result in a more anatomical keyboard layout. Comparing the angles between each neighboring pair of keys in all rows

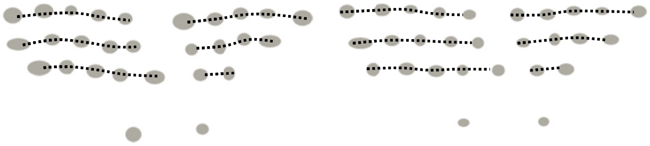

(a) ***asterisk feedback, no keyboard*** **(b)** ***asterisk feedback, visible keyboard***

Figure 7. Centroids of key hits with contour ellipses of two standard deviations in *x* and *y* directions. Bezier curves fitted to the centroids illustrate curvature. Each participant's data is translated to the same midpoint between F and J keys. (*N* = 20)

(except for the split between hands at T-Y, G-H, and B-N) provides an estimate of the deviation from a straight line for that row, where an angle of 0º would be perfectly horizontal. The mean absolute angle between each pair of keys in the *no keyboard* condition was 9.9º compared to the *visible keyboard* condition at 5.4º (significant difference: $t_{19} = 7.35$, $p < .001$). Thus, a curved keyboard design should best support touch-typing.

Distance between hands is greatest in the 'no keyboard' condition. We computed the average distance between the rightmost keys of the left hand and the leftmost keys of the right hand (between pairs T-Y, G-H, and B-N). On average, there was 1.12" of space between the hands in the *visible keyboard* condition ($SD = 0.19$") and 1.41" of space with *no keyboard* ($SD = 0.42$"), a difference that was significant ($t_{19} = 4.17$, $p = .001$). The mean distance in the *visible keyboard* condition is more than the 0.9" of space between visual key centers ($t_{19} = 5.30$, $p < .001$). This result shows that users are most comfortable typing with a gap between their hands. Even with a visible keyboard, the underlying key press model may need to take this gap into account.

Hit Point Deviations per Key

Where the keyboard shape analysis examined centroids of key hits, a more detailed examination of the spread of hits per key allows us to identify individual keys that may benefit from an increase in size. We first calculated the standard deviation of hits for each key per participant in *x*- and *y*-directions. We then grouped the 26 letter keys by finger and row, since previous research has shown these factors can affect error rates [4]. For example, the Q, A, Z, and P keys were grouped as *little* finger. For each of *x*- and *y*-direction standard deviation, we ran a repeated measures ANOVA with the following within-subjects factors: *typing input* (*no keyboard* vs. *visible keyboard*), *row* (bottom, middle, top), and *finger* (little, ring, middle, index).

Overall, hit point deviations were greatest in the 'no keyboard' condition. There was a main effect of *keyboard* for both *x*- and *y*-directions (*x*-direction: $F_{1,19} = 10.77$, $p = .004$; *y*-direction: $F_{1,19} = 39.28$, $p < .001$). These results reflect the pattern evident in Figure 6, that there was a smaller spread of hits for each key when participants were given visual constraints compared to when they were not.

The little finger resulted in the greatest horizontal spread of hits. There was a main effect of *finger* on *x*-direction deviation ($F_{1.6,31.1}$=5.79, $p = .011$). Pairwise comparisons showed the keys assigned to the little finger had significantly greater *x*-direction deviation than the ring ($p = .033$) and middle fingers ($p = .024$), while comparison to the index finger was only a trend ($p = .075$). No other significant main or interaction effects were found on *x*-direction deviation. This finding suggests that keys pressed with the little finger should be widest.

Bottom row keys, especially with 'no keyboard', resulted in the greatest vertical spread of hits. All main and interaction effects were significant for *y*-direction deviation, so we examine the highest-order effect in detail: a three-way interaction of *keyboard* × *finger* × *row* ($F_{2.8,34.3} = 5.70$, $p = .002$). Significant pairwise comparisons (at $p < .05$) showed differences were stronger with *no keyboard*: the bottom row resulted in greater *y*-direction deviation than the middle and top rows for the little and index fingers, and just the top row for the ring finger. With the *visible keyboard*, this pattern held, but was only significant with the index finger. These results suggest that making the keys in the bottom row taller may improve accuracy. For completeness, the other significant effects were: *finger* ($F_{3,57} = 17.88$, $p < .001$), *row* ($F_{2,38} = 16.11$, $p < .001$), *keyboard* × *finger* ($F_{3,57} = 4.65$, $p = .006$), *keyboard* × *row* ($F_{1.3,23.7} = 5.59$, $p = .021$), and *finger* × *row* ($F_{1.3,3.7} = 6.96$, $p = < .001$).

Key-to-Hand Mappings

The keyboard shape analyses showed that participants were most comfortable with a gap between their hands. Although the most obvious split would be based on the standard key-to-hand mapping for touch-typing (i.e., left hand: T, G, B and keys to the left; right hand: Y, H, N and keys to the right), an analysis of actual key-to-hand mappings indicates that even skilled typists have idiosyncrasies in this respect.

Spacebar use is predominantly by only one thumb. Almost all participants used only one thumb for the spacebar (right: 14 participants; left: 4 participants), replicating previous results with a wearable keyboard [6]. There was no relationship between handedness and thumb choice.

Middle keys were often shared between hands. Based on the 12 participants for whom we logged complete hand contour data, we checked within which hand each key press occurred. Some participants used the opposite hand or alternated hands for the B, H, and Y keys (e.g., right hand for B key). Excluding potential noise from mislabeled data when there were few (< 5) key presses from a participant using the opposite hand, we saw the left hand accounted for 16.7% of Y presses (3 participants) and 5.2% of H presses (1 participant), while the right hand accounted for 11.4% of B presses (4 participants). In 0.15% of cases, the center of the touch point was offset such that it fell between the hand contours; half of such cases occurred with H.

Key Press Classification

The analyses presented above provide insight into how keyboard layout and key size may be improved to support touch-typing patterns on a flat surface. In this section, we assess the reliability of key hit locations to evaluate how accurate the modified designs could be. We perform simple distance-based classification of key presses, both within a participant, and between a participant and the group's average. Again, the *no keyboard* condition offers the closest representation to eyes-free touch-typing.

User-dependent key press classification is highest with the visible keyboard, yet still 90% with no keyboard. Using 10-fold cross-validation, we calculated the centroid of key presses for each training subset of the data, and classified the remaining key presses based on the closest centroid. Key strikes were relatively reliable for both keyboard conditions, although more so for the *visible keyboard*. Mean classification accuracy was 90.0% (*SD* = 5.4%) for *no keyboard* and 96.7% (*SD* = 3.7%) for the *visible keyboard*. The range was from 75.5% to 97.5% for *no keyboard* and 87.0% to 99.8% for the *visible keyboard*. On the whole, subjects were consistent within themselves, repeatedly hitting the same places for the same keys.

Classification accuracy per letter is shown in Figure 8. Keys in the bottom row were the most difficult to classify with *no keyboard*, which follows our earlier result that bottom row keys had the highest *y*-direction deviation. The U, I, J, and K keys were also relatively difficult to classify in both keyboard conditions. This finding requires more investigation, but may be related to hand displacement that could occur after a backspace gesture with the right hand.

Incorrect classifications most often occurred as adjacent keys in the same row. Based on the classification results, we created confusion matrices for each condition. Overall, most incorrect classifications were adjacent keys, vertically or horizontally, with most occurring in the same row (67.7% with the *visible keyboard* and 54.8% with *no keyboard*). Examining whether misclassifications occurred to the left, in the same column, or to the right, results were more evenly split: only 20.7% and 30.9% occurred in the same column for the *visible keyboard* and *no keyboard* conditions, respectively. With the *visible keyboard* condition, the most frequent misclassifications were J→K (4.0% of instances), L→K (3.4%), O→I (2.6%) and U→I (2.6%). The most frequent misclassifications in the *no keyboard* condition were B→V (7.6%) and V→B (7.5%). The predominance of same-row errors in the data, suggests that participants found it easier to reliably hit keys in the vertical direction than horizontally.

User-independent classification lowers accuracy by 19.5% in the 'no keyboard' condition. We calculated the average of key centroids for all users, translated so they were centered on the midpoint between the F and J keys. We then classified each user's input based on the average centroids. The classification accuracy dropped from the user-dependent analysis to, on average, 93.1% (*SD* = 8.1%) for the *visible keyboard* condition and 70.5% (*SD* = 15.2%) for the *no keyboard* condition. The large drop in accuracy from the user-dependent classification for the *no keyboard* condition indicates that some degree of personalization is necessary for eyes-free touch typing.

Summary of Major Results

In the *unrestricted* typing condition, users were able to achieve fast typing speeds (59 WPM) when they assumed their input was correct. Extra finger touches were largely constrained to between typing sequences, while non-finger touches were spatially segmented from the finger touches. With *asterisk feedback*, the keyboard shape in the *no keyboard* condition was more curved and had a greater gap between hands than the *visible keyboard* condition. There was a larger spread of hits per key with *no keyboard* than the *visible keyboard*. However, key presses could still be classified with 90% accuracy in the *no keyboard* condition using a simple user-dependent classification method. Finally, some keys were pressed with both hands, and bottom-row keys and keys associated with the little fingers had the greatest hit point deviations.

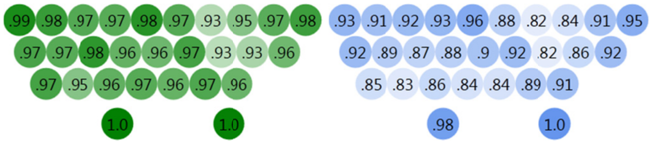

(a) ***asterisk feedback, visible keyboard*** **(b)** ***asterisk feedback, no keyboard***

Figure 8. User-dependent classification accuracy per key, showing difficulty with bottom row with *no keyboard*. (N = 20)

DISCUSSION

Our goals were to identify typing patterns that could improve ten-finger, flat-surface keyboards, and to investigate the feasibility of touch-typing on flat surfaces.

Design Recommendations

Inactive palm area. To filter spurious hand and arm presses, three columns of space may be used, centered horizontally at the midpoint between thumbs and extending downward. The left and right columns are inactive, while the middle column may be used for input (including the spacebar). Using 4" columns, this simple heuristic would eliminate 95% of non-finger touches in the *asterisk feedback* conditions, and impact only 0.2% of finger touches.

Spacebar width. Spacebar hits in our visualizations suggest the spacebar should be narrower and taller than a traditional spacebar. Reducing the size of keys where possible has the potential to reduce spurious touch points. Thus, we recommend a spacebar extending from the middle of C to the middle of M for a keyboard of the size tested here.

Keyboard curvature. Rather than arranging each row of keys straight across, the *asterisk feedback, no keyboard* condition demonstrates that rows should be arched, more representative of relative finger lengths and reach.

Space between hands. Ten-finger keyboards should allow for a gap between hands. When no constraints were imposed on hand placement (*asterisk feedback, no keyboard*), the average distance between the centers of the rightmost keys on the left side of the keyboard and the leftmost keys on the right side of the keyboard was 1.41".

Relative key sizes. The keys assigned to the little finger and the bottom row keys had greater *x*- and *y*-direction deviation in key press locations, respectively, than other keys. These keys should be relatively larger.

Key-to-hand mappings. Keys in the middle of the keyboard, especially B, H, and Y, are sometimes typed with either hand. This finding is relevant to split keyboard designs, or designs that place keys relative to hand locations. Crossover was especially common with the B and Y keys, which could potentially be placed on both sides of a split keyboard.

In general, a more ergonomic layout should improve input accuracy over a rectangular one. In the *asterisk feedback, visible keyboard* condition, only 81.5% of key presses occurred within key bounds. However, key press classification based on the emergent key centroids for all users was almost 12% higher.

We did not explore different visual designs, but the visual affordances of the keyboard would affect typing patterns. Many of the recommendations listed here could be implemented with or without a visual affordance. For example, allowing for space between the hands could mean the underlying keyboard model adjusts key centers away from the middle of the keyboard, but it does not necessarily mean that a visual gap must appear. Future work should explore what the best visual affordance, if any, will be for each design recommendation, and how user behavior changes with respect to visual changes in the keyboard.

Towards Touch-Typing on Flat Surfaces

The goal of this study was not only to identify design recommendations for current whole-hand touch screen keyboards, but also to explore the feasibility of eyes-free touch-typing on a flat surface. In the *unrestricted* typing condition, where participants were not aware of input errors, mean typing speed was 59 WPM. This number is indicative of speeds that novice users could achieve with an ideal ten-finger touch screen keyboard, and performance should improve with use. The *unrestricted* condition was slower than the physical keyboard, which we speculate may be due to differences in the mechanics of the two setups (i.e., flat surface vs. raised keyboard) and to previous negative experiences with touch screens, which could have made some participants initially hesitant. Again, with more practice these effects should decrease.

Key press classifications from the *asterisk feedback, no keyboard* condition also point to the potential for touch-typing on a flat surface. With no visual constraints, classification accuracy with a simple user-dependent model was 90%. While 90% is hardly perfect, it could be improved through more sophisticated classification schemes. Language modeling would also improve performance further (e.g., [7,14,30]). Word-level correction approaches should be particularly effective here: such approaches require clean segmentation between words, and the spacebar was the easiest key to classify.

Personalization will most likely be a key element of any flat surface keyboard that allows for touch-typing. We observed many individual differences in terms of spacing between hands, size and shape of key press distributions, and key-to-hand mappings. Underscoring these findings, user-dependent key press classification was about 20% more accurate than user-independent classification for the *asterisk feedback, no keyboard* condition. This disparity suggests that if we want to allow users to type without frequently looking at the keyboard (as in that condition), the underlying model will need to adapt to each user.

Limitations

The conditions studied here provided either no feedback or masked feedback to users. While this decision was necessary to achieve our goals, providing users with unmasked text output would certainly impact behavior, allowing users to adapt their typing patterns to create more accurate output if necessary. We plan to explore methods to improve input accuracy, such as intelligently identifying spurious touch points or using a pressure sensitive surface. Further study is needed on how close experienced users will come to achieving the ideal speeds seen in the *unrestricted* condition. We predict that in real typing tasks users will achieve speeds somewhere between the *asterisk feedback* conditions and the *unrestricted* condition.

Our participants only included expert touch-typists, which is a critical user group to study if the goal is to design the most efficient text input methods possible. Although the *no keyboard* conditions are not directly applicable to novice typists, the resulting design recommendations may still improve performance for those users; for example, the observed differences in hit point deviations per key may be reflective of basic human motor performance. Future work will need to confirm the degree to which the *visible keyboard* findings also apply to novice typists.

Finally, we required users to place their hands consistently at the start of each trial to reduce noise from potential hand drift over the course of the study. Without this requirement, we would expect a decrease in the reliability of key press locations. Detecting the location of the user's hands and adjusting the keyboard if needed may be a useful approach for mitigating this issue. Language modeling could also be used to offset this projected decrease in accuracy.

CONCLUSION

We have investigated the unconstrained typing patterns of 20 expert typists on a flat surface. Our results demonstrate that typing patterns differ when users are provided with a visual keyboard compared to no visual affordance, yet key press locations remain relatively reliable within an individual. Design recommendations emerging from this study should improve the effectiveness of static touch screen keyboard designs. But our vision is to design keyboards that will allow users to touch-type on a flat surface. The results presented here should encourage researchers to pursue this goal, and indicate that an effective solution will require an element of personalization. Future work should apply the design recommendations here, and investigate the potential to which touch-typing on flat surfaces can be achieved.

ACKNOWLEDGMENTS

The authors thank Andrew D. Wilson and Jon Froehlich for their help and feedback. This work was supported in part by the National Science Foundation under grant IIS-0811063, NSERC, Microsoft Surface, and Intel Labs.

REFERENCES

1. Al Faraj, K., Mojahid, M., and Vigouroux, N. 2009. BigKey: A Virtual Keyboard for Mobile Devices. *Proc. HCII*, 3-10.
2. Barrett, J. 1994. Performance effects of reduced proprioceptive feedback on touch typists and casual users in a typing task. *Behaviour & Information Tech. 13 (6)*, 373-381.
3. Benko, H., Morris, M. R., Brush, A.J.B., Wilson, A.D. 2009. Insights on Interactive Tabletops: A Survey of Researchers and Developers. *Microsoft Research Technical Report MSR-TR-2009-22*. March, 2009.
4. Cooper. E. (ed.). 1983. *Cognitive Aspects of Skilled Typewriting*. Springer-Verlag, New York, NY.
5. Go, K., and Endo, Y. 2007. CATKey: customizable and adaptable touchscreen keyboard with bubble cursor-like visual feedback. *Proc. HCII*, 493-496.
6. Goldstein, M., Book, R., Alsiö, G., and Tessa, S. 1999. Non-keyboard QWERTY touch typing: a portable input interface for the mobile user. *Proc. CHI'99*, 32-39.
7. Goodman, J., Venolia, G., Steury, K., and Parker, C. 2002. Language modeling for soft keyboards. *Proc. AAAI'02*, 419-424.
8. Grudin, J.T. 1984. Error patterns in skilled and novice transcription typing. In *Cognitive Aspects of Skilled Typewriting*, W. E. Cooper (ed.). New York: Springer-Verlag, 121-143.
9. Gunawardana, A., Paek, T., and Meek, C. 2010. Usability guided key-target resizing for soft keyboards. *Proc. IUI '10*, 111-118.
10. Hartmann, B., Morris, M. R., Benko, H., and Wilson, A. D. 2009. Augmenting interactive tables with mice & keyboards. *Proc. UIST '09*, 149-152.
11. Himberg, J., Häkkilä, J., Kangas, P., and Mäntyjärvi, J. 2003. On-line personalization of a touch screen based keyboard. *Proc. IUI'03*, 77-84.
12. Hinrichs, U., Hancock, M., Collins, C., Carpendale, S. 2007. Examination of text-entry methods for tabletop displays. *Proc. Tabletop 2007*, 105-112.
13. Kristensson, P. and Zhai, S. 2004. SHARK^2: a large vocabulary shorthand writing system for pen-based computers. *Proc. UIST '04*, 43-52.
14. Kristensson, P. and Zhai, S. 2005. Relaxing stylus typing precision by geometric pattern matching. *Proc. IUI '05*, 151-158.
15. Langendorf, D.J.: Textware solution's Fitaly keyboard v1.0 easing the burden of keyboard input. *WinCELair Review*, February 1998.
16. Lankford, C. 2000. Effective eye-gaze input into Windows. *Proc. ETRA '00*, 23-27.
17. McAdam, C., Brewster, S. 2009. Distal tactile feedback for text entry on tabletop computers. *Proc. BCS-HCI'09*, 504-511.
18. MacKenzie, I.S 2002. A note on calculating text entry speed. http://www.yorku.ca/mack/RN-TextEntrySpeed.html.
19. MacKenzie, I.S., Soukoreff, R. W. 2003. Phrase sets for evaluating text entry techniques. *Ext. Abstr. CHI '03*, 754-755.
20. MacKenzie, I.S., Zhang, S. X. 1999. The design and evaluation of a high-performance soft keyboard. *Proc. CHI '99*, 25-31.
21. MacKenzie, I.S., Zhang, S. X. 2001. An empirical investigation of the novice experience with soft keyboards. *Behaviour & Information Technology*, *20*, 411-418.
22. Morris, M.R., Lombardo, J., Wigdor, D. 2010. WeSearch: Supporting Collaborative Search and Sensemaking on a Tabletop Display. *Proc. CSCW 2010*, 401-410.
23. Rashid, D. R., Smith, N. A. 2008. Relative keyboard input system. *Proc. IUI'08*, 397-400.
24. Roeber, H., Bacus, J., Tomasi, C. 2003. Typing in thin air: the Canesta projection keyboard—a new method of interaction with electronic devices. *Proc. CHI '03*, 712-713.
25. Ryall, K., Forlines, C., Shen, C., Ringel Morris, M., Everitt, K. 2006. Experiences with and Observations of Direct-Touch Tabletops. *Proc. Tabletop 2006*, 89-96.
26. Salthouse, T.A. 1986. Perceptual, cognitive, and motoric aspects of transcription typing. *Psych. Bulletin 99(3)*, 303-319.
27. Sears, A., Revis, D., Swatski, J., Crittenden, R., Shneiderman, B. 1993. Investigating touchscreen typing: the effect of keyboard size on typing speed. *Behavour & Information Technology 12(1)*, 17-22.
28. Shein, F., Hamann, G., Brownlow, N., Treviranus, J., Milner, M. and Parnes, P. 1991. WiViK: A visual keyboard for Windows 3.0. *Proc. RESNA '91*, 160-162.
29. Soukoreff, R. W. and MacKenzie, I. S. 2001. Measuring errors in text entry tasks: an application of the Levenshtein string distance statistic. *Proc. CHI '01*, 319-320.
30. Tinwala, H., & MacKenzie, I. S. (in press). Eyes-free text entry with error correction on touchscreen mobile devices. *Proc. NordiCHI 2010* (to appear).
31. Weiss, M., Wagner, J., Jansen, Y., Jennings, R., Khoshabeh, R., Hollan, J. D., and Borchers, J. 2009. SLAP widgets: bridging the gap between virtual and physical controls on tabletops. *Proc. CHI '09*, 481-490
32. Wigdor, D., Penn, G., Ryall, K., Esenther, A., Shen, C. 2007. Living with a Tabletop: Analysis and Observations of Long Term Office Use of a Multi-Touch Table. *Proc. Tabletop 2007*, 60-67.
33. Wobbrock, J.O., Morris, M.R., Wilson, A.D. 2009. User-defined gestures for surface computing. *Proc. CHI '09*, 1083-1092.
34. Wobbrock, J.O. and Myers, B.A. 2006. Analyzing the input stream for character-level errors in unconstrained text entry evaluations. *ACM TOCHI 13(4)*, 458-489.
35. Zhai, S., M. Hunter, and B.A. Smith. 2000. The Metropolis Keyboard—an exploration of quantitative techniques for virtual keyboard design. *Proc. UIST 2000*, 119-218.
36. Zhai, S. and Kristensson, P. 2003. Shorthand writing on stylus keyboard. *Proc. CHI '03*, 97-104.
37. Zhai, S., Smith, B.A., Hunter, M. 2002. Performance Optimization of Virtual Keyboards. *Human-Computer Interaction 17(2&3)*, 229 - 269.
38. Zhai, S., Sue, A., Accot, J. 2002. Movement model, hits distribution and learning in virtual keyboarding. *Proc. CHI '02*, 17-24.