A Comparative Evaluation of Spatial Targeting Behaviour Patterns for Finger and Stylus Tapping on Mobile Touchscreen Devices

DANIEL BUSCHEK, JULIA KINSHOFER, and FLORIAN ALT, LMU Munich, Germany

Models of 2D targeting error patterns have been applied as a valuable computational tool for analysing finger touch behaviour on mobile devices, improving touch accuracy and inferring context. However, their use in stylus input is yet unexplored. This paper presents the first empirical study and analyses of such models for tapping with a stylus. In a user study ($N = 28$), we collected targeting data on a smartphone, both for stationary use (sitting) and walking. We compare targeting patterns between index finger input and three stylus variations – two stylus widths and nib types as well as the addition of a hover cursor. Our analyses reveal that stylus targeting patterns are user-specific, and that offset models improve stylus tapping accuracy, but less so than for finger touch. Input method has a stronger influence on targeting patterns than mobility, and stylus width is more influential than the hover cursor. Stylus models improve finger accuracy as well, but not vice versa. The extent of the stylus accuracy advantage compared to the finger depends on screen location and mobility. We also discuss patterns related to mobility and gliding of the stylus on the screen. We conclude with implications for target sizes and offset model applications.

CCS Concepts: • **Human-centered computing** → **Empirical studies in ubiquitous and mobile computing**;

Additional Key Words and Phrases: Stylus input, offset model, Gaussian Process regression, computational interaction

1 INTRODUCTION

Computational methods and tools have gained importance and attention in HCI research on mobile touch interfaces and interactions, for example for interface optimisation [50] and adaptation [11, 39, 62], layouting of keyboards [21, 42, 51, 72] and homescreens [60], prediction of touch input [27, 44, 48] and occluded areas [64], as well as detection of cognitive errors [45]. Common to this work are models which help to anticipate and recognise what (future) user input might likely look like. In other words, these computational methods and models enable touch devices to capture and utilise *expectations*. One such expectation concerns future touch targeting behaviour, in particular errors (offsets) relative to the intended target. To reduce such offsets and improve touch accuracy, the outlined computational toolbox emerging from recent HCI research proposed *touch offset models* [26, 67]. These models predict the user's actually intended touch location from a given sensed touch location.

Related work [10, 14, 15, 66, 67] has examined the underlying 2D touch-to-target offset patterns on mobile touchscreen devices for finger input, showing that finger touch accuracy can be significantly improved by modelling a user's individual targeting pattern with such offset models. Moreover, these models can be employed to

Authors' addresses: Daniel Buschek, Julia Kinshofer, Florian Alt, LMU Munich, Amalienstrasse 17, 80333, Munich, Germany.

computationally inform GUI design, for example to estimate touch locations for GUI elements and layouts [12, 14], and to infer context information such as hand (left vs right [15]) and finger (thumb vs index [10]).

Overall, accurate targeting (and modelling it) is challenging since this behaviour is influenced by many factors, such as occlusion [6, 63], grasp and reachability [7], finger orientation [30], perception and properties of the implement [2, 31], and movement [22]. Based on related work [10, 26, 67], offset models present a promising approach to capturing the resulting patterns and to improve accuracy.

However, 2D offset patterns have not been analysed for stylus input yet and it is not clear if offset models improve stylus accuracy. The use and investigation of styli on (mobile) devices has a long history, including early tablets [41], PDAs [43, 54, 70], and recent touchscreen devices [1, 2, 19, 34, 61]. Styli also enjoy commercial success (e.g. Samsung's Galaxy Note series[1], Apple Pencil[2]). Since styli will thus likely continue to be a part of future mobile touchscreen devices and since offset patterns are influenced by many factors (e.g. hand posture and target type [14], mobility [46]), we argue that it is important to evaluate them for input methods beyond finger touch. Related work supports this view: Tu et al. [61] recently highlighted the importance of fundamental comparative evaluations of finger and stylus, for example to inform future design of interfaces and recognition algorithms. Similarly, a detailed study on understanding stylus accuracy concluded with a call for further exploration of "the diverse and complex nature of accuracy" [2]. Sharing these views, we argue that it is important to extend the computational toolbox for targeted input on mobile touch devices to include offset models for styli as well.

To this end, we contribute: 1) the first study and analyses of 2D targeting offset patterns for stylus tapping; 2) a detailed comparison of targeting behaviour and offset modelling between stylus and finger tapping for both stationary and mobile use; and 3) a dataset of stylus tapping behaviour.

2 RELATED WORK

We relate our investigation to work on analysing, comparing, and modelling (mobile) stylus and finger input.

2.1 Finger Touch Behaviour and Spatial Models

Finger touch input on mobile devices suffers from several sources of inaccuracies, such as the finger's softness and its occlusion of targets [6, 63], limited thumb reach [7], varying pitch, roll, and yaw [30], body movement [23], encumbrance [49], and the individual use of visual finger features [31]. These factors lead to varying touch locations for a given target location.

For the specific case of mobile touch keyboards, this is often taken into account with probabilistic models of the keys [25, 72]. More generally, the spatial distribution of touches around keys and other targets was modelled with a single Gaussian [5, 65] or a combination of two Gaussians [8]. Further influences on mobile touch targeting behaviour include hand posture [5, 14] and grip [44], as well as GUI target shape, location, and size [14, 38], the target's proximity to screen edges [4], and also environmental temperature [56].

Another view focusses on offsets between intended target and actual touch locations. Related work modelled these offsets based on touch targeting data to predict and correct future touches for more accurate touch interaction. Henze et al. [26] derived polynomial touch offset functions from large datasets in-the-wild. Weir et al. [67] highlighted the user-specific and non-linear aspects of the underlying targeting behaviour patterns, proposing non-linear models. They further investigated techniques to reduce the training set size [66]. User-specificity of offset patterns was also explored for its biometric value [14], and compared to device-specific influences [15]. Moreover, Musić et al. [46, 47] examined influences of walking. Finally, offset models were employed to infer hand postures [10, 15], to analyse influences of hand size [10], and to examine and predict touch behaviour on mobile websites [12]. In this paper, we adopt these offset models to analyse stylus input.

Further patterns related to finger touch were studied as well, for example regarding device tilt and task completion time [17]. These patterns were examined in a grid across the screen. The same work also used differential patterns (i.e. patterns of differences between conditions) to compare input techniques. We also analyse differential patterns to compare stylus and finger input. In contrast to the related work, we employ regression models to avoid binning data into a fixed grid.

In summary, offset models and underlying spatial targeting patterns remain unexplored for stylus input. We thus apply and evaluate offset models to analyse stylus targeting patterns and to improve stylus accuracy.

2.2 Stylus Input Behaviour and Spatial Models

The use of a stylus with interactive surfaces has a long history, from the light pen's use in Sutherland's seminal work on *Sketchpad* [59] and early handwriting recognition systems [20] to today's smartphones and tablets.

Here, we focus on sources of inaccuracy for stylus input. This is a known problem of practical importance, that keeps users from adopting styli for very precise input tasks [2], along with other issues such as unintended touch and device latency [1, 16, 49]. Stylus inaccuracy also resulted in the first explorations of combining targeting models and language models for mobile text entry [24], a line of research leading to the foundations of gesture typing [36, 37], and continuing to this day for modern touchscreen keyboards [51, 72].

Factors influencing accuracy include the presence of visual feedback and the size of the stylus tip (nib) [2]. Recommendations for stylus length and width were derived as well [53], also for different age groups [54]. In our study, we thus also examine targeting patterns for different stylus variations: a thin stylus with and without hover cursor, and a "thick" stylus. By analysing the (user-specific) spatial structure of stylus targeting behaviour across the screen for the first time, we follow the related work's call for detailed exploration of accuracy with more attention to the user [2].

A main topic of existing models of stylus input is palm rejection [3, 57] – discriminating unintended palm touches from intended input. This is an important practical problem that can result in ergonomic issues [16]. Hinckley et al. [28] proposed another discriminative model to distinguish different ways of holding the stylus in combined finger touch and stylus interaction on tablets. In contrast, Vogel et al. [64] proposed a model for hand occlusion in direct stylus input on a tablet. They later employed this model to realise occlusion-aware touch interfaces with stylus input [62].

With the exception of the latter, these models are discriminative: They improve stylus applications by distinguishing patterns, but they do not capture how users produce those. In contrast, other work probabilistically modelled targeted stylus tap distributions [24], similar to later work on finger input [5, 8, 65]. In this paper, we extend the toolbox of fundamental stylus models by investigating targeting patterns and offset models for stylus input. These offset models are generative in the sense that they can "imagine" a distribution of likely intended stylus input locations for a sensed location. Moreover, they can be applied inversely to simulate input for given target locations, useful for computational analysis of GUIs [12, 14].

2.3 Combinations and Comparisons of Finger and Stylus Input

Brandl et al. [9] and Hinckley et al. [28, 29] proposed interaction techniques that combine finger and stylus input on a touchscreen, following earlier work that included such combinations [68, 71].

However, most related work compared these and other input methods: Early work by Mack and Lang [40] compared mouse, stylus and finger touch on a PC and found the finger to be slightly slower than the stylus. They also pointed out accuracy problems for the finger. MacKenzie et al. [41] found that the stylus outperformed mouse and trackballs during pointing. A similar study by Kabbash et al. [35] also took into account differences between dominant and non-dominant hand.

Several studies show that the stylus is more accurate than the finger: Lee and Zhai [38] evaluated finger and stylus tapping with on-screen buttons on mobile devices. Their results are in favour of the stylus for both error rate and speed. Holzinger et al. [33] compared finger and stylus performance in selection tasks on tablet PCs used by staff in a hospital in sitting, standing, and walking contexts. The stylus was significantly faster and more accurate than the index finger. Cockburn et al. [19] compared target acquisition performance of mouse, stylus and finger for tapping and (radial) dragging on a tablet PC. The finger was overall fast but inaccurate.

Beyond speed and accuracy there is still little work on the differences and similarities between stylus and finger. Recent related work [61] also made this observation and noted that this leads to the transfer of interfaces and interactions from stylus to finger input without further investigation (e.g. gesture keyboards [73]).

Our results confirm that the stylus is more accurate than the index finger. However, in contrast to all previous work we compare spatial targeting behaviour not as averages but as detailed 2D patterns across the screen. In this way, we paint a more refined picture, as motivated above. This includes pattern analyses of horizontal and vertical offsets, user-specific targeting, offset shifts while walking, and gliding (tap down to up distances). In summary, our results reveal interesting spatial structure behind the previously observed average differences.

Tu et al. [61] compared finger and stylus stroke gestures on mobile devices, both while sitting and walking. They found that stylus and finger strokes differ in several features, like gesture size ratio, yet are similar in others. Most interestingly for our investigation, they found that the stylus had better accuracy than the index finger while sitting, but comparable performance while walking. We study tapping, not strokes, and also found that walking reduces accuracy differences between stylus and index finger. We reveal that while walking the right index finger "catches up" with stylus accuracy in particular along the lower right edge of the screen, demonstrating additional insights gained through detailed analyses of spatial patterns.

3 USER STUDY

3.1 Background: Modelling Spatial Targeting Patterns

We give a short introduction to touch offset modelling. In the following, we use the term "touch" both for finger touch as well as a tap with a stylus. Formally, an offset model is a function $f(t)$ which maps a touch location $t = (t_x, t_y)$ to an offset $o = (o_x, o_y)$. Adding this offset to the touch yields a corrected touch location $t' = (t_x + o_x, t_y + o_y)$. This function can take on different shapes and is learned from data via regression. Related work proposed both linear models (polynomials [26], quadratic basis functions [10, 15]) and non-linear models (Gaussian Processes [10, 14, 67], Relevance Vector Machines [14, 66]). Some models predict not only one offset for each touch location, but rather a whole distribution, to account for uncertainty [10, 14, 15, 67]. In this paper, we mainly use Gaussian Process models [10, 14, 67], since they are flexible and do not assume a particular class of functions in advance, in contrast to linear regression models. We refer the reader to this related work for more details on these models and their implementation.

3.2 Study Design

We use a within-subject design with two independent variables. The first one is *input method* with the four levels *hover* (thin stylus with hover cursor), *thin* (the same stylus without a hover cursor), *thick* (a thicker stylus), and *finger* (the right index finger). The second independent variable is *mobility* with the two levels *sitting* and *walking*. These conditions were motivated by related work, which found differences between such styli [2], and also influences of mobility on offset patterns [46]. Hence, to provide a comprehensive picture, we decided to include mobility in our investigation of stylus offset patterns as well.

For each combination of these variables, we collected touches for targets shown at 300 different screen locations on a grid. For our analyses, we recorded the following information per touch: timestamp touch down/up, touch down/up locations (x, y), and target location (x, y).

Fig. 1. The two styli used in the study.

3.3   Participants

We recruited 28 participants (11 female, mean age 25 years, range 18-38 years) via our university's mailing list and social media. All but one participant were right handed. Participants received a 15 € gift card for an online shop. Most people had little prior experience with using a stylus on a mobile device.

3.4   Apparatus

We used a Samsung Galaxy Note 4 smartphone, running a custom Android app that displayed an empty white screen with a single cross-hair as a target. Tapping triggered the display of the next target, as explained in the procedure below. The app also displayed the current target count next to the required number of target acquisitions. After completing this number of targets, the app displayed instructions for the next task (e.g. whether to use the finger or stylus).

We used two styli (Figure 1): The Samsung S-Pen (top) is included in the Samsung Galaxy Note 4 and has a plastic nib with a diameter of 1.2 mm. The other one was a thicker "Powery" stylus (bottom) with a rubber nib and a diameter of 4.5 mm. For the Samsung S-Pen, the Samsung Galaxy Note 4 can display a hover cursor. This is a small grey circle that indicates the stylus location (x, y) when the stylus nib is hovering closely above the screen. Related work showed that a hover cursor improves stylus input accuracy [2]. However, this hover cursor system cannot take into account parallax effects based on the relative location of head, stylus and screen.

3.5   Procedure

The study took place in a quiet room and hallway in a university building in two sessions per participant. Each session lasted for about 30 to 45 minutes. We first explained the study and its tasks. Participants then signed a consent form, followed by the main part of the study, explained below.

Tasks: Each participant completed the same eight targeting tasks per session – sitting and walking with all four input methods. The order of tasks was counterbalanced with a Latin Square as far as possible ($N = 28$ but 8×8 square, i.e. last iteration of the square not complete).

Targeting: In each task, participants had to hit 300 cross-hair targets, shown one at a time by our app. These targets were shown in random order, taken from a grid to ensure that targeting was observed for the whole screen. Each touch triggered the display of the next target. Participants received no feedback on "hits" or other targeting performance measures.

Mobility: Participants sat on a chair for the four tasks in the sitting condition. In the four walking tasks, they were instructed to walk up and down a straight corridor of about 30 m length without any obstacles at their normal pace. Some people sometimes slowed down considerably or stopped walking at all. In these cases, the instructor reminded them to again walk at their normal pace.

Questionnaire: After completing the eight tasks, the first session concluded with a questionnaire on demographics and perception of the tasks. Participants returned for a second session after five days to two weeks (most returned after a week, but we arranged alternative dates for people who were busy exactly a week later). In the second sessions, they again completed the eight targeting tasks described above.

3.6 Limitations

Our study has a few limitations: The two styli differ slightly in shape and size, and in particular in both nib diameter as well as nib material. Hence, differences between the two cannot be directly associated with just one factor. Most people in our sample were novice stylus users. Moreover, we used one phone model and examined an abstract targeting task with cross-hairs, not realistic GUI elements. Similarly, we studied a walking task on a predefined course. Mobility in real life is likely more dynamic than our setting. On the other hand, the lab study allowed us to control such external influences to measure fundamental targeting behaviour. Similar to finger touch [26], future work could also study stylus targeting patterns beyond the lab over a longer time period. Furthermore, our study sample has a limited age range. It will be interesting to compare these results, for example, to those obtained for elderly people. Moreover, with one exception, our sample is limited to right handed people. We thus do not investigate the influence of the dominant hand on offset patterns.

4 RESULTS

4.1 Preprocessing

We normalised all touch and target locations to the zero to one range, as in related work [10, 14, 15, 67]. Sometimes, participants hit the screen by accident, for example tapping two times in a row. This creates outlier touches that are far away from the displayed target. To not distort the patterns, we remove these outliers as in related work [10, 14]: Outliers are touches which are further away from the corresponding target than that task's mean offset length plus three standard deviations. We removed 2713 touches as outliers in this way (2.02 % of the data).

4.2 Analysis Overview

To unclutter the following report, we state general measuring, modelling, and testing procedures here once: To test for significance, we conduct repeated measures ANOVAs, using Greenhouse-Geisser correction for violations of sphericity, and Bonferroni correction to account for multiple tests. Significance is reported at $p < 0.05$.

In the following, we measure accuracy as root-mean-square-error (RMSE) of the touch-to-target distances (i.e. offset vector lengths). Hence, a lower RMSE indicates higher accuracy.

All models are trained per user (i.e. on data from one user – "user-specific"), following related work [10, 14, 15, 67]. We train Gaussian Process (GP) models. GPs require us to specify a small set of hyperparameters. We informed the value for the γ hyperparameter (see [10]) with a grid search using data from one randomly chosen user, resulting in $\gamma = 0.4$. This approach was also used in related work [67]; it avoids "peaking into the testing data", apart from that one user. For the other hyperparameters, we use the values from related work [10, 14].

Finally, we exclude the left-handed participant for all pattern analyses, since left-handed patterns are clearly different from right-handed ones [10, 15].

4.3 Targeting Accuracy and Absolute Improvements

As a first analysis of touch accuracy and improvements with touch offset models, we compared absolute RMSEs on raw touches to RMSEs on corrected touches (factor *model* in the following analyses). The models used for correction were trained per user on data from one task in session one and tested on data from the same user and task in session two.

Table 1 (left) shows the descriptive statistics. Figure 2 visualises these results. Comparing RMSEs, we found significant effects of *mobility* ($F(1, 27) = 73.222$, $p < 0.001$, $\eta^2 = 0.731$), *input method* ($F(2.012, 54.325) = 84.674$, $p < 0.001$, $\eta^2 = 0.758$) and *model* ($F(1, 27) = 74.739$, $p < 0.001$, $\eta^2 = 0.735$). All interactions were also significant: *mobility × input method* ($F(2.411, 65.099) = 4.548$, $p < 0.009$, $\eta^2 = 0.144$), *mobility × model* ($F(1, 27) = 13.594$, $p < 0.001$, $\eta^2 = 0.335$), *input method × model* ($F(1.683, 45.431) = 36.134$, $p < 0.001$, $\eta^2 = 0.572$), and *mobility × input method × model* ($F(2.104, 56.803) = 13.088$, $p < 0.001$, $\eta^2 = 0.326$).

Targeting accuracy (RMSE), sitting				Targeting accuracy (RMSE), walking				Accuracy improvement (% RMSE)			
input method	model	M	SD	input method	model	M	SD	mobility	input method	M	SD
finger	raw	0.035	0.009	finger	raw	0.040	0.007	sitting	finger	17.070	9.859
	GP	0.029	0.009		GP	0.037	0.007		hover	9.262	6.843
hover	raw	0.021	0.012	hover	raw	0.031	0.013		thin	8.585	8.536
	GP	0.020	0.012		GP	0.030	0.012		thick	4.747	7.090
thin	raw	0.022	0.010	thin	raw	0.032	0.012	walking	finger	7.517	5.435
	GP	0.020	0.009		GP	0.030	0.011		hover	5.906	4.958
thick	raw	0.025	0.011	thick	raw	0.032	0.010		thin	5.331	4.395
	GP	0.024	0.011		GP	0.031	0.010		thick	3.496	5.228

Table 1. Left and centre: Descriptive statistics for targeting accuracy (RMSE), both for raw touches and touches corrected with GP offset models. Right: Descriptive statistics for relative improvements (% improvement with models in RMSE).
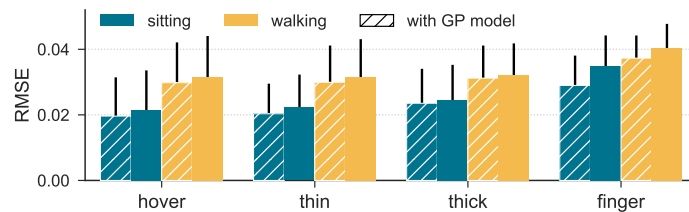

Fig. 2. Touch targeting accuracy for the different study tasks. The figure shows root-mean-squared-errors (RMSEs) of touch-to-target distances per task in the second session. Lower values indicate higher accuracy. We show RMSEs both for the raw touch locations as well as touch locations corrected with a user- and task-specific Gaussian Process (GP) offset model, trained on the corresponding data from the first session, collected a week earlier.

To break down these interactions, contrasts were performed comparing walking to sitting, the model's corrected touches to the raw ones, and all styli to the finger. This revealed significant interactions when comparing walking to sitting both for hover compared to finger ($F(1, 27) = 6.829, \eta^2 = 0.202$) and thin compared to finger ($F(1, 27) = 5.134, \eta^2 = 0.160$). The interaction graphs showed that the negative influence of walking on accuracy is thus significantly stronger for the thin stylus compared to the finger, but not for the thick stylus.

Moreover, the analyses showed significant interactions when comparing RMSE with the model's corrected touches to the raw touches for all styli compared to the finger (hover: $F(1, 27) = 35.876, \eta^2 = 0.571$; thin: $F(1, 27) = 37.094, \eta^2 = 0.579$; thick: $F(1, 27) = 54.901, \eta^2 = 0.670$). The interaction plots showed that the accuracy improvement achieved with the model is thus significantly stronger for the finger than for the styli (compare descrease in RMSE between raw and GP between finger and styli in Table 1 left/centre).

The analyses also revealed significant interactions when comparing walking to sitting for the model's corrected touches compared to the raw touches ($F(1, 27) = 13.594, \eta^2 = 0.335$). The negative influence of walking on accuracy is significantly stronger for the model compared to the raw touches. However, looking at the values in detail shows that this is only the case for the finger (RMSE +0.008 vs +0.005 – compare increase in RMSE between sitting and walking for raw vs GP in Table 1 left/centre). This also reflects in the significant three-way interactions when comparing walking to sitting, for the model's corrected touches compared to the raw touches, for the styli compared to the finger (hover: $F(1, 27) = 19.124, \eta^2 = 0.415$; thin: $F(1, 27) = 14.110, \eta^2 = 0.343$; thick: $F(1, 27) = 22.678, \eta^2 = 0.456$).

In summary, styli are significantly more accurate than the finger. The thin stylus with hover is the most accurate stylus, followed by thin and thick stylus without cursor. The results also show for all tasks that offset models trained for one user and task significantly reduce RMSE for that user and task. We repeated this analysis with linear offset models [10, 15] and found only negligible differences. To avoid redundancy, the following analyses thus continue to report only on the GP models.

			testing task (2nd session)							
			sitting				walking			
			finger	hover	thin	thick	finger	hover	thin	thick
training task (1st session)	sitting	finger	17.07	-29.25	-20.42	-19.10	5.93	-10.93	-9.94	-8.80
		hover	6.71	9.26	7.30	1.31	3.58	4.65	4.10	0.86
		thin	8.82	7.19	8.59	1.31	4.64	4.14	4.11	0.39
		thick	8.32	1.16	3.84	4.75	4.46	2.51	2.40	2.73
	walking	finger	13.30	-22.75	-14.42	-14.08	7.52	-6.67	-5.51	-4.59
		hover	5.45	3.19	5.66	-1.38	3.86	5.91	5.89	0.84
		thin	7.50	1.60	5.88	-0.69	4.90	4.85	5.33	0.95
		thick	5.03	-5.65	-0.46	1.95	3.88	0.08	0.28	3.50

Table 2. Relative accuracy improvements in percent, achieved with GP offset models trained and tested across tasks.

4.4 Relative Accuracy Improvements

We also analysed *relative* improvements, computed as $100(\text{RMSE}_{raw} - \text{RMSE}model)/\text{RMSE}_{raw}$. Hence, "perfect" improvement of 100 % would mean that the model corrected all touches to be exactly at their intended target locations. We evaluated these improvements by training models on data from session one and testing on the same task in session two (as in the previous section).

Table 1 (right) shows the descriptive statistics. In summary, the mean improvement for all finger conditions was 12.29 %, and the mean improvement for all stylus conditions was 6.22 %. These improvements are in line with related work for the finger [10, 14, 26, 67]. We found significant effects of *mobility* ($F(1, 27) = 28.442, p < 0.001, \eta^2 = 0.513$) and *input method* ($F(2.417, 65.260) = 18.545, p < 0.001, \eta^2 = 0.407$) on improvements. The interaction effect was significant as well ($F(2.783, 75.142) = 9.304, p < 0.001, \eta^2 = 0.256$).

Contrasts revealed significant differences between all styli and the finger (hover: $F(1, 27) = 13.827, \eta^2 = 0.339$; thin: $F(1, 27) = 17.515, \eta^2 = 0.393$; thick: $F(1, 27) = 39.326, \eta^2 = 0.593$). Looking at the values (Table 1 right) shows that the relative improvement is thus significantly higher for the finger than the styli. Contrasts also showed significant interactions when comparing walking to sitting for all styli compared to the finger (hover: $F(1, 27) = 35.876, \eta^2 = 0.571$; thin: $F(1, 27) = 37.094, \eta^2 = 0.579$; thick: $F(1, 27) = 54.901, \eta^2 = 0.670$). The interaction graph showed that the negative influence of walking on accuracy improvements is significantly stronger for the finger compared to the styli. In summary, while the finger significantly benefits relatively more from offset model corrections, the styli retain a significantly larger proportion of their models' improvements when users start walking compared to sitting. Overall, offset models relatively reduce RMSEs significantly more for sitting than walking, and more so for finger input than for the thin stylus with hover cursor, followed by thin and thick stylus without cursor.

4.5 Relative Accuracy Improvements Across Tasks

For cross-task analysis, we evaluate improvements by training models on data from one task in session one and testing them on another task in session two. Table 2 summarises the results, with within-task improvements along the diagonal. These results show that the finger benefits from any model, even one trained on stylus data in a different mobility context. On the other hand, the finger model can only improve finger input, also in a different mobility context, but not stylus input. Moreover, models trained on sitting data applied to walking input are overall more robust than vice versa: Sitting finger models decrease accuracy less when applied to walking styli than walking finger models applied to sitting styli. Moreover, stylus models trained on sitting data do not decrease accuracy for walking (positive values in top right quadrant). In contrast, stylus models trained on walking data decrease accuracy for sitting in several cases (negative values in bottom left quadrant). These results are useful to inform enrolment practices (see discussion).
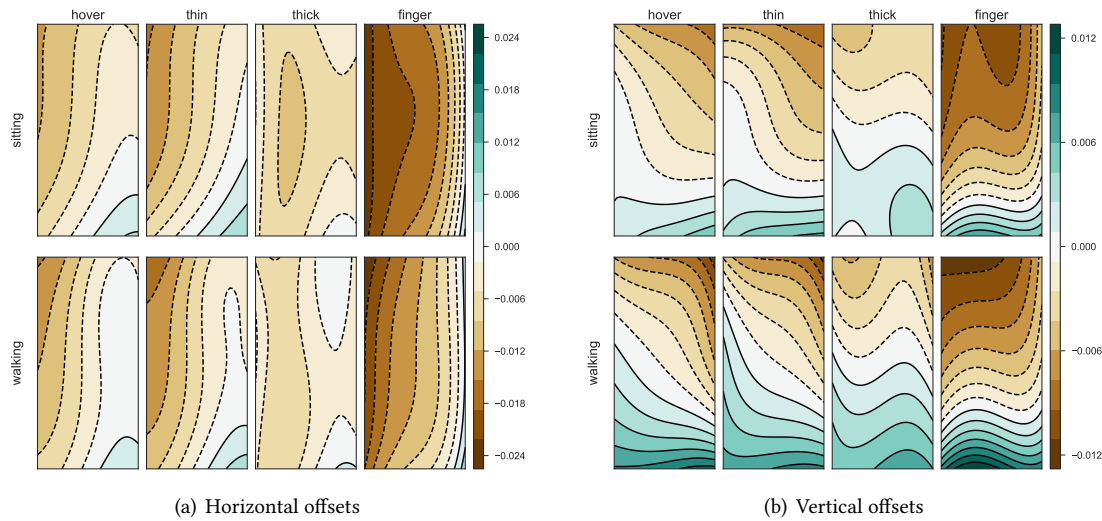
(a) Horizontal offsets

(b) Vertical offsets

Fig. 3. Touch targeting offset patterns per study task, obtained by averaging the predictions of user- and task-specific models across participants: *(a)* Horizontal (x) offsets, from "touched too far to the left relative to the target" (positive values, teal colour, full lines) to "touched too far to the right" (negative values, brown, dotted lines). *(b)* Vertical (y) offsets, from "touched above the target" (positive values, teal, full lines) to "touched below the target" (negative values, brown, dotted).

4.6 Offset Patterns

Figure 3 shows the mean patterns per task. These patterns were created by computing offsets for a grid of screen locations with all user-specific models, then averaging across users per location. In general, the pattern for the index finger is considerably different from the three stylus variations.

4.6.1 Horizontal Offsets. Looking at the horizontal patterns in Figure 3(a), we notice that the index finger pattern shows a clear tendency to touch too far to the right, relative to the targets. This is likely the result of the targeting angle and reaching movement with the right index finger, matching similar results in related work [10]. Moreover, related work suggests that the relative location of head, device, and finger influences targeting [31].

In contrast, the thin stylus patterns (with and without hover cursor) show a diagonal structure: At its extremes, users tend to hit too far to the right near the top left corner of the screen, and too far to the left near the bottom right corner. This suggests that users tend to avoid getting close to the screen corners with the stylus (cf. [4]).

Finally, the pattern for the thick stylus is relatively balanced, in between finger and thin stylus, with less right-shift than the finger, but also less diagonal structure than the thin stylus. Overall, these observations hold for both sitting and walking contexts.

4.6.2 Vertical Offsets. For all input methods, vertical patterns (i.e. y-offsets) show a common basic structure: users tend to hit below targets near the top of the screen, and above targets near the bottom. Again, this indicates that users try to avoid hitting the screen edges.

For the finger, and in particular while walking, vertical offsets are largest near the two left corners, whereas the hover/thin stylus patterns show largest vertical offsets (i.e. darkest regions in Figure 3(b)) near the right corners. Similar to horizontal offsets, we explain this difference between stylus and finger with the reaching angle and movement required to bring the index finger tip down to the screen: Larger offsets near the left corners suggest that people tend to minimise their reaching movements, which for right hand input get longer towards the left. Azenkot and Zhai reported a similar observation [5].
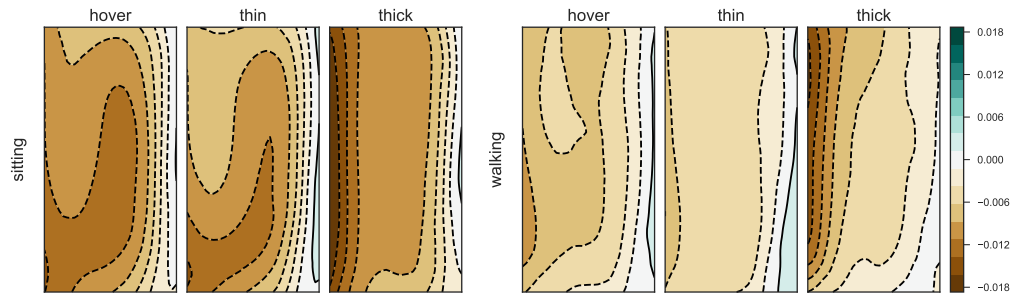
Fig. 4. Difference models comparing accuracy between stylus and index finger input. This figure plots difference in offset vector lengths between task-specific models trained on stylus and finger data. Brown indicates areas in which the stylus is more accurate, while the index finger is more accurate in teal screen regions. While styli are overall more accurate than the index finger, this advantage is not consistent across the whole screen, especially while walking.

4.7 Stylus vs Finger Patterns

This section examines differences between styli and the index finger in more detail. We compare accuracy patterns across the screen and evaluate how characteristic and consistent input method-specific targeting behaviour is.

4.7.1 Accuracy Differences. We modelled accuracy differences between finger and stylus across the screen: For each stylus task and user, we computed the offset length predicted at each screen location and subtracted the corresponding offset length for the finger. Figure 4 shows these difference models averaged over all users.

Overall, the stylus is more accurate than the finger, matching results from related work [19, 33, 38]. However, our pattern analyses reveal that this difference depends on the tap location: The right index finger is almost as accurate as the stylus very close to the right edge of the screen, and worse at other locations. We explain this with a greater aversion to hit the device edge with a hard stylus compared to a soft finger, since recent related work observed a reluctance towards touching the outside of the screen [4]. Along the left edge, though, this is overshadowed by the finger's lower accuracy (see previous section), resulting in the described left-to-right pattern. Moreover, this right edge region in favour of the finger is more prominent during walking than sitting. This is in line with results in recent related work on mobile touchscreen gesture drawing, where walking also reduced the differences in accuracy between finger and stylus input [61]. For tapping, our results show that this overall reduction of differences relates to certain screen regions.

4.7.2 Recognition and Consistency. Another way of analysing differences in offset patterns between stylus and finger is by employing them for input method recognition: Intuitively, if patterns are different, they should help to recognise from tap data whether a user taps with a stylus or a finger.

To this end, we evaluate each user's offsets observed in the second session under the predictive distributions from both a model trained on finger data and one trained on stylus data from that user's first session. We accumulate the resulting input method probabilities over time (i.e. over multiple taps). For more details on this approach to classification with offset models and its implementation we refer the reader to related work [10, 14].

We observed accuracies of 80-90 % correct stylus-vs-finger predictions after about 60 touches for all three stylus variations, both while sitting and walking (TNR about 85-93 %, FNR about 10-25 %). This result demonstrates that offset patterns are *characteristic* regarding the input method (stylus vs finger), with *consistent* differences across sessions a week apart.
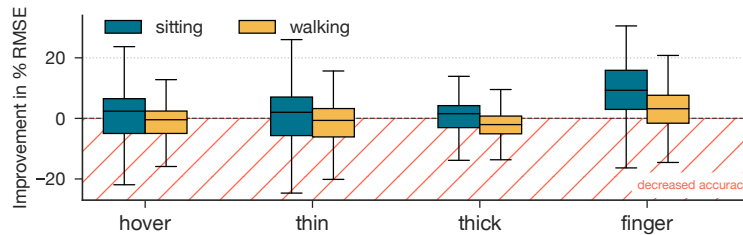
Fig. 5. Offset model applications across users. This boxplot shows the distributions of relative accuracy improvements obtained by training and testing task-specific models on data from different users, for all pairs of users. "Negative improvement" means that the model results in worse accuracy (red area). These results show that targeting behaviour is highly user-specific.

4.8 User-Specificity of Targeting Behaviour Patterns

Related work found that touch offset patterns are user-specific (see e.g. [67]). This can be explained by the many influences on finger targeting that may vary from user to user, such as hand size [10], grasp and finger orientation [30] and visual targeting strategy [31]. User-related factors were also identified as influences for stylus accuracy [1, 2]. Hence, we expect that styli produce user-specific patterns as well, motivating this analysis.

To examine user-specificity, we evaluated applications of offset models across users: We trained models per task on data from one user u_i in session one and evaluated them on data from another user u_j in session two. We repeated this for all pairs of users (u_i, u_j), excluding the left-handed user. For each task, this results in $27 \times 26 = 702$ RMSE improvement values. Figure 5 summarises these improvements as boxplots.

The figure shows that overall models did not improve a user's accuracy when trained on a different user, as evident by many "negative improvements" (i.e. the model's corrections made accuracy worse). In particular, no stylus input method could be reliably improved with cross-user models. In contrast, finger touch while sitting was the only task for which cross-user models could improve accuracy for most users. In summary, these results thus suggest that 2D targeting patterns and related offset corrections are even more user-specific for stylus input than for finger touch.

4.9 Influence of Mobility on Targeting Patterns

To investigate the influence of mobility on offset patterns, we computed difference models comparing offsets between sitting and walking contexts: For each input method and user, we computed the x/y offsets predicted at each screen location while sitting and subtracted the corresponding x/y offsets for the walking condition, respectively. Figure 6 visualises these difference models averaged over all users.

Overall, our results show that walking affects horizontal offsets for the thin stylus in a more complex way than for the finger. We explain this as follows: The (relative) movements of the holding and interacting hand while walking lead to the general change in offsets for both stylus and finger. However, compared to the finger and thick stylus, the thin stylus glides further on the screen while walking, causing the differences in the offset shifts between finger and thin stylus. We next present further analyses on gliding which support this explanation.

Glide is measured between touch down and up location. The mean glide distances for sitting were 0.004 for finger, 0.010 for hover, 0.011 for thin, and 0.003 for thick. While walking, mean glide distances were higher: 0.007 for finger, 0.017 for hover, 0.018 for thin, and 0.006 for thick. Glide distance was significantly influenced by *mobility* ($F(1, 27) = 146.09, p < 0.001, \eta^2 = 0.844$) and *input method* ($F(1.281, 34.593) = 86.97, p < 0.001, \eta^2 = 0.763$). The interaction effect was also significant ($F(1.887, 50.936) = 29.53, p < 0.01, \eta^2 = 0.522$).

Contrasts revealed significant differences between the thin stylus and the finger (hover: $F(1, 27) = 70.149, \eta^2 = 0.722$; thin: $F(1, 27) = 123.085$). Looking at the values showed that the glide is thus significantly higher for the
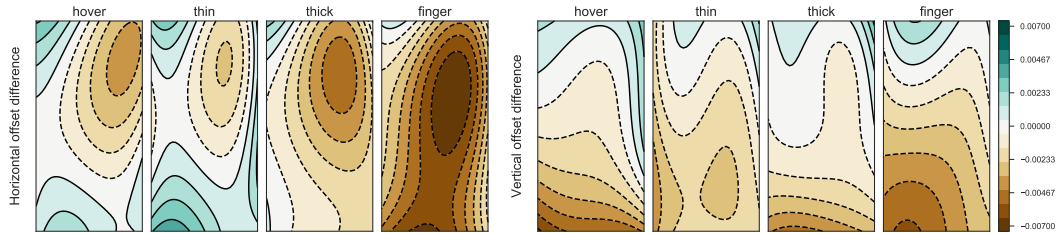
Fig. 6. Difference models comparing offsets between sitting and walking contexts, averaged over all participants. This figure plots offset differences between task-specific models trained on sitting and walking data. Brown indicates areas in which x/y offsets shifted to the left/top, respectively, when walking compared to sitting. Complementary, teal shows where walking shifted x/y offsets towards the right/bottom.
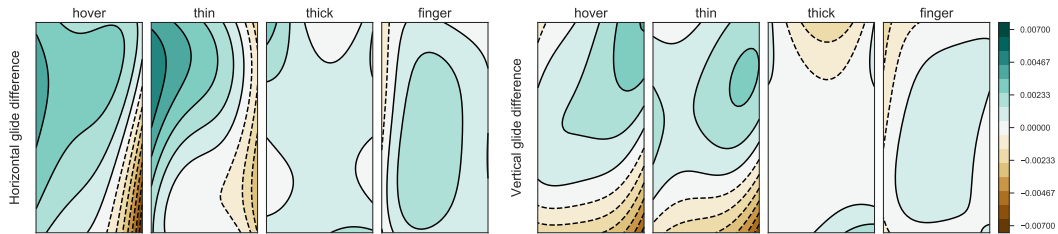

Fig. 7. Difference models comparing glides (i.e. distances between tap down and up) between sitting and walking, averaged over all participants. This figure plots difference between task-specific glide models trained on sitting and walking data. Brown indicates areas in which x/y glides shifted to the right/bottom, respectively, when walking compared to sitting. Teal shows where walking shifted x/y glides towards the left/top. In particular, walking resulted in glides shifting towards the screen edges for the thin stylus with and without hover cursor. Thick stylus and index finger were less affected.

thin stylus than the finger, but not for the thick stylus. Contrasts also showed significant interactions when comparing walking to sitting for the thin styli compared to the finger (hover: $F(1, 27) = 48.428, \eta^2 = 0.642$; thin: $F(1, 27) = 58.203, \eta^2 = 683$), but not for the thick stylus. In summary, the thick stylus and the finger form a "short glide" group, and the increase in gliding due to walking is significantly stronger for the thin stylus.

Overall, the thin stylus thus glided the most, whereas the thick one and the finger had more "grip". This difference is explained by the different materials (plastic vs rubber, see Figure 1). These quantitative results fit to the subjective reports of a lack of friction between a similar thin stylus and the screen in related work [1].

Figure 7 visualises the glide differences between sitting and walking. The strongest differences (darkest colours) are observed for the thin stylus. They form something akin to a "watershed" along which walking shifts glides either towards the left/top or the right/bottom. Location and shape of this border presumably are a result of the relative movement of the device (and the supporting left hand) and the interacting right hand while walking.

4.10 Subjective Feedback

After completing the tasks, participants provided subjective feedback via a short questionnaire and free comments. We summarise this feedback as follows: Overall, 64 % (18 people) chose the finger as their preferred input method, followed by 21 % (6 people) for the thin stylus with hover cursor. In contrast, the thin stylus without cursor and the thick stylus were each preferred by two people only. Most people used the finger more often than the stylus in their everyday interactions and thus likely preferred their more familiar method here. In free comments, the thin pen was criticised by some as too thin and lightweight to allow for a good grip. Other comments indicated the well-known problem of accidentally hitting the screen with the hand when using styli.

We also asked participants whether the hover cursor had helped them in the targeting tasks: 50 % (14 people) rated it favourably (4 and 5 on 5-point Likert scale). In contrast, 35 % (10 people) did not find it helpful (ratings 1 and 2). Comments showed that some people found the cursor distracting or that they felt pressured to be more precise when using it, since the cursor revealed inaccurate targeting.

Finally, regarding mobility, most people recognised and commented on the increased difficulty of targeting while walking, compared to sitting. Some said that the hover cursor did not help them while walking and that they thus found finger and thick pen more useful here.

5 SUMMARY AND DISCUSSION
We summarise and discuss the key insights from our study and analyses.

5.1 Insights into Targeting Behaviour Patterns and Modelling
We expect these findings and lessons learned to be useful for applications of offset models and future evaluations of both finger touch and stylus input:

Offset models improve stylus tapping accuracy both while sitting and walking: Our results show that offset models previously known from finger touch can be adopted for stylus tapping as well. They significantly improve stylus tapping accuracy by 4.74 % to 9.26 % when sitting, and by 3.50 % to 5.91 % while walking.

Stylus offset models also improve finger touch, but not vice versa: Offset models trained on data from stylus tapping also improve accuracy for finger touch. In contrast, finger models do not improve stylus input.

One explanation is overall higher accuracy with styli: Since offsets are shorter for styli, the models learn to predict shorter offsets, compared to models trained on finger data. In that sense, styli models are more conservative, meaning they shift touches less far. This "carefulness" is beneficial for transferring models across tasks, where they might not fit exactly (e.g. from stylus to finger input). In contrast, the finger models' larger touch corrections can shift touches too far when applied to stylus data. This is particularly the case in those regions where the finger is a lot less accurate than the stylus (left screen border, likely due to people minimising reaching movements [5]).

Fingers benefit relatively more from offset corrections than styli: Offset models relatively improve accuracy significantly more for finger input than for the thin stylus with hover cursor, followed by thin and thick stylus without cursor. Note that the order is not simply defined by the baseline accuracy (e.g. the thick stylus is more accurate than the finger, but benefits relatively less from the models' corrections). This refines the result from related work on finger touch that offset models show relatively larger improvement the less accurate the user is to begin with [10, 15]: Our results show that improvement also depends on the implement (finger vs stylus) and its properties (nib size, material).

Since offset models are trained on previous touches, they better predict offsets if behaviour stays consistent [14]. Hence, the larger improvements for finger than styli likely also result from less consistent patterns with styli than the finger between the two study sessions. For example, users might have held the styli (slightly) differently in the two sessions (for grip variations see e.g. [1, 58]). In contrast, there is no such additional "tool grip" variation for finger input. This is supported by related work [14], which found that index finger patterns are less characteristic and consistent than thumb targeting; index finger input uses a different kinematic chain (i.e. involving the whole arm), in contrast to holding and operating the phone with one hand and its thumb. It thus seems plausible that adding a stylus further affects consistency and in turn decreases accuracy benefits of offset modelling.

Input method has a stronger influence on targeting offset patterns than mobility: We found significant effects on accuracy for both input method and mobility, as well as their interaction. The underlying 2D offset patterns differed between finger and stylus, more so than between sitting and walking per method.

We explain this as follows: While walking affects accuracy and gliding, so does using a stylus instead of the finger. However, walking appears to shift offsets less systematically than changing the input method, which affects several systematic influences on targeting (e.g. nib size / finger affects occlusion [2], and likely visual target alignment, cf. [31]).

Stylus type has a stronger influence on targeting offset patterns than a hover cursor: Overall, 2D offset patterns differed more between thin and thick stylus than between thin stylus with and without cursor. Related work [2] concluded the opposite, yet they compared average accuracy, not 2D targeting patterns, and our studies are different in that our styli differ in both nib size as well as material (plastic vs rubber). Future work could investigate 2D offset patterns for the different plastic nibs from the related work [2].

We explain our result as follows: While the hover cursor (slightly) increases accuracy (i.e. shortening offsets), it does not change the overall offset pattern. In contrast, the impact of a different nib size and material may change the patterns for several reasons (also known from finger input, see related work section), such as target occlusion, different stylus grip, and visual alignment with the target.

Stylus targeting offset patterns are highly user-specific: Related work has shown that finger touch offsets are user-specific [14, 67]. Based on our analyses with offset models applied across users, we conclude that stylus targeting is highly user-specific as well. In our study, cross-user applications of offset models only worked rather reliably for finger input. This indicates that stylus tapping is even more individual than targeting with the index finger. As mentioned before, styli introduce a source of variability related to the way that users hold them [1, 58]. This is likely to vary with user-specific factors such as hand size and experience [2], explaining individual patterns.

A thin plastic stylus nib glides more and in a more complex pattern than finger and thick stylus: Overall, the thin stylus with the plastic nib glided about 150 % more on the screen than finger and thick stylus with rubber nib. Comparing the underlying 2D gliding patterns revealed that the thin stylus has a more complex bipolar "gliding dynamic", in particular while walking. Moreover, walking increases gliding with the thin stylus significantly more than for the finger and thick stylus. These results shed light on subjective reports of a lack of stylus friction in related work [1]. Overall, the gliding differences can be explained by the different friction of the nib materials (and the finger's skin).

The accuracy advantage of the stylus compared to the finger depends on screen location and mobility: We confirm the general result from previous work (e.g. [19, 33, 38, 61]) that the stylus is more accurate than the index finger. However, our analyses of 2D offset patterns reveal a more refined picture: the extent of this advantage depends on the screen area. The right index finger "catches up" with stylus accuracy near the right screen edge, in particular while walking. This supports and refines the results of Tu et al. [61], who observed that finger and stylus had comparable performance while walking, but not while sitting.

Overall, our refined pattern of the stylus accuracy advantage is explained by several related observations: Users minimise reaching movements [5], resulting in both larger horizontal offsets on the left side and larger vertical offsets near the left screen corners. Fittingly, one participant commented that these targets required moving his arm more. The stylus can help with extra reach. In addition, offset directions indicate that people tend to avoid hitting the device borders (cf. [4]), possibly more so for the hard stylus than the soft finger. This is particularly the case while on the move, as evident from the comparison of patterns between sitting and walking: Presumably, people account for walking when trying to avoid the device borders. This explanation seems plausible considering the finding from related work that people consider their gait cycle when typing [47].

5.2 Comparison to Other Approaches

Besides offset models, there exist other approaches to help with targeting. For example, Vogel and Baudisch proposed the *Shift* [63] technique. It shows a callout with a copy of the occluded area when touching small targets. An earlier related concept is the *Offset Cursor* [52], which places a cursor at a fixed offset above the finger. Roudaut et al. [55] proposed further ideas: *TapTap* uses one tap to magnify an area of interest and a second one for selection. *MagStick* uses dragging to bring up a "stick" that extends thumb reach and snaps to targets.

One main advantage of offset models is that they are interface-agnostic, meaning that they only need the touch location and no further information about the GUI. However, related work showed that offsets are influenced by the target shape and size [14], so depending on the application some target-related information might still be beneficial. Moreover, offsets are user-specific, thus training per user is recommendable.

More generally, dedicated interaction techniques offer more control for the user. In contrast, if a touch is corrected by an offset model the user cannot further influence this. A comparison of performance and perceived control between offset models and such techniques in different targeting tasks presents an interesting direction for future work. Furthermore, combinations of offset predictions with targeting concepts could also be explored. For example, one drawback of the *Offset Cursor* is its fixed offset. If placed above the finger, it is difficult to hit targets near the bottom of the screen. Using a dynamic offset informed by an offset model might help here.

5.3 Implications for Design and Applications

Implications for GUI target placement: Related work found that the target's proximity to screen edges [4] can negatively impact on finger touch targeting performance. Our analyses of targeting patterns support this finding – largest offsets occurred near screen edges and corners. Based on our patterns and the related work we thus recommend to ensure a gap between screen edges and GUI elements.

Implications for GUI target size and adaptation: Figure 8 shows recommended diameters of buttons in millimetres for both finger and stylus input, averaged across all participants and both sitting and walking contexts. These sizes were derived from our data and (inverse) offset models by assuming that users target the centres of circular buttons, sized such that the mean offset plus three standard deviations would still be a hit.

Importantly, these patterns connect varying results and recommendations from previous research and industry: For example, several studies on finger touch input recommended sizes of 20 mm and found no further benefits beyond that (see [18]). Similarly, our finger pattern plateaus at 20 mm and reveals that this is only relevant in a very limited screen region. Fittingly, Lee and Zhai [38] found sizes of 10 mm acceptable for stylus input on a smartphone's lower half of the screen (i.e. keyboard area); the finger resulted in more errors there. Their device had a comparable physical size to ours. In line with this, our patterns 1) recommend smaller target sizes, even below 10 mm, in similar regions (bottom half / bottom right quadrant); and 2) they also capture that the stylus can deal with smaller targets there. The minimum target size for finger touch suggested by Holz and Baudisch [31] (4.3 mm) also falls into the value range of the finger's most accurate region in our pattern. Moreover, Google's Android design guidelines[3] recommend button sizes of 7 mm to 10 mm. According to our patterns, this is more than large enough for the stylus across the whole screen, and large enough for the right index finger in most areas apart from the top left corner.

Regarding applications, related work highlighted the need for fundamental comparisons of stylus and finger to support adaptations of UIs and interaction techniques from one to the other [61]. In this view, designers could use our derived patterns to help inform decisions on whether and how to relocate and resize GUI elements to adapt their interface from finger to stylus input or vice versa; or to accommodate both.

---

[3]https://material.io/guidelines/layout/metrics-keylines.html#metrics-keylines-touch-target-size
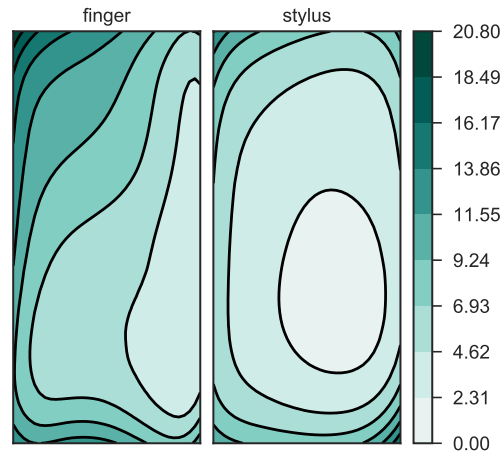
Fig. 8. Guidelines for GUI target size. Colour shows the recommended (minimum) diameter of circular buttons in millimetres (values: read upper end per colour). The patterns are based on data across all participants and both sitting and walking.

A disclaimer is necessary: We highlight that these patterns were derived from limited data from a single device model and should thus not be misunderstood as a final global design rule. Nevertheless, they clearly indicate that spatial structure is relevant when examining input accuracy on handheld touch devices and may help explain and bring together some of the varying results from the literature, as outlined above.

Finally, our patterns also reveal that the most accurate areas – affording very small buttons – are biased towards the right for right-handed input. In contrast to one-handed thumb use [7], reachability is not an issue with index finger and stylus; this right-shift is thus rather explained by the observation that users try to minimise finger travel distance, as pointed out by Azenkot and Zhai [5]. Hence, our pattern analyses clearly support the idea of hand posture-adaptive GUI widgets, not just for one-handed (thumb) use [11], but also for index fingers and styli.

Implications for enrolment practices: A mobile device using both finger and stylus input might initially ask users to tap on some targets to train a user-specific offset model [66]. We show: To keep this enrolment task short, it can ask for stylus input only, since stylus models also improve finger input. In contrast, it should not ask for finger input only, since finger models do not improve stylus accuracy. Moreover, it is possible to conduct an enrolment task only while sitting, as models trained on sitting data improve accuracy while walking as well. Since improved accuracy indicates that models fit the user's individual behaviour well, these findings might also help to inform enrolments for biometric user authentication/identification with offset models (see [14]).

Implications for touch biometric systems: Research on adaptive mobile touch UIs [23, 72] as well as touch biometrics [13] addressed challenges caused by hand posture-specific variations in touch behaviour (e.g. typing with thumb vs index finger). Similarly, recent related work [14] on touch targeting biometrics concluded that systems should assess hand postures and GUI layouts. Extending this view, mixed finger and stylus input should be considered as well, since their targeting patterns are clearly different, as revealed in our analyses here. Following ideas from the literature, this could be done by integrating over both finger and stylus patterns in a probabilistic model [13]; or with a hierarchical "backoff" model that in addition to thumb(s) and index finger(s) also includes a sub-model for stylus input [14, 72]. Such sub-models could use the offset models discussed here, since our results show that they capture targeting behaviour specific to both user and input method.

6 CONCLUSION

Computational methods and models enable touch devices to capture and utilise expectations about user behaviour, useful to analyse and improve interfaces and interactions. One such expectation concerns touch targeting errors (offsets) relative to the intended target. Recent HCI research [10, 14, 15, 66, 67] proposed touch offset models which predict the user's actually intended touch location from a given sensed touch location. Besides their use in analysing behaviour patterns, these models can significantly improve finger touch accuracy.

However, touch offset patterns and models have not been analysed and evaluated for stylus input yet. To close this gap, this paper reported on the first user study ($N = 28$) and analyses of 2D targeting offset patterns for stylus tapping. We compared targeting behaviour on a smartphone between three stylus variations and the index finger, both while sitting and walking. Moreover, to the best of our knowledge, we are the first to model 2D gliding patterns, and to derive GUI target sizing patterns from offset models.

Our results reveal that offset models significantly improve stylus tapping accuracy, but less so than for finger input. Evaluating patterns in detail we find that: 1) stylus targeting behaviour is highly user-specific; 2) input method has a stronger influence on targeting patterns than mobility; 3) stylus width is more influential than the hover cursor; 4) stylus models improve finger accuracy as well, but not vice versa; 4) a thin stylus with plastic nib glides more and has more complex gliding patterns than the index finger and a thick stylus; and 5) the extent of the well-known average accuracy advantage of the stylus compared to the finger varies for different screen areas, in particular while walking.

Overall, we conclude that offset models are a useful computational tool for 1) researchers to analyse stylus input and to compare targeting behaviour between stylus and finger; and for 2) practitioners to design GUIs with spatial accuracy patterns in mind and to improve input accuracy.

7 FUTURE WORK

Future work could examine targeting patterns for further styli and devices, for example miniature styli on smart watches [69]. Follow-up studies could also extend the set of participants, for example to elderly people or users with motor impairments. Such user groups might display rather different offset patterns than the ones observed here (see e.g. [18]). Other GUI elements could be examined as well, as finger touch offset patterns are influenced by target shapes and sizes [14]. Moreover, offset models could be extended to make use of other sensors beyond the touchscreen, such as inertial sensors attached to the stylus.

Beyond tapping, styli are often used in continuous tasks such as drawing or handwriting. The usefulness of offset models here is not clear yet: These models might be applied to improve accuracy of the initial location. There might also be some value in shifting the location continuously (e.g. to cancel effects on recognition related to different tilting in different screen regions), yet this requires careful further investigation.

Targeting behaviour patterns could also be examined and compared between styli and finger for the full capacitive sensor data, not just the touch location derived by the system (see e.g. [32]).

8 PROJECT RESOURCES

Data and pattern plots as well as modelling scripts and tools are available at:
   http://www.medien.ifi.lmu.de/stylus-patterns/

9 ACKNOWLEDGEMENTS

REFERENCES

[1] Michelle Annett, Fraser Anderson, Walter F. Bischof, and Anoop Gupta. 2014. The Pen is Mightier: Understanding Stylus Behaviour While Inking on Tablets. In *Proceedings of Graphics Interface 2014 (GI '14)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 193–200. http://dl.acm.org/citation.cfm?id=2619648.2619680

[2] Michelle Annett and Walter F. Bischof. 2015. Hands, Hover, and Nibs: Understanding Stylus Accuracy on Tablets. In *Proceedings of the 41st Graphics Interface Conference (GI '15)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 203–210. http://dl.acm.org/citation.cfm?id=2788890.2788926

[3] Michelle Annett, Anoop Gupta, and Walter F. Bischof. 2014. Exploring and Understanding Unintended Touch During Direct Pen Interaction. *ACM Trans. Comput.-Hum. Interact.* 21, 5, Article 28 (Nov. 2014), 39 pages. https://doi.org/10.1145/2674915

[4] Daniel Avrahami. 2015. The Effect of Edge Targets on Touch Performance. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1837–1846. https://doi.org/10.1145/2702123.2702439

[5] Shiri Azenkot and Shumin Zhai. 2012. Touch Behavior with Different Postures on Soft Smartphone Keyboards. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '12)*. ACM, New York, NY, USA, 251–260. https://doi.org/10.1145/2371574.2371612

[6] Patrick Baudisch and Gerry Chu. 2009. Back-of-device Interaction Allows Creating Very Small Touch Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1923–1932. https://doi.org/10.1145/1518701.1518995

[7] Joanna Bergstrom-Lehtovirta and Antti Oulasvirta. 2014. Modeling the Functional Area of the Thumb on Mobile Touchscreen Surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1991–2000. https://doi.org/10.1145/2556288.2557354

[8] Xiaojun Bi, Yang Li, and Shumin Zhai. 2013. FFitts Law: Modeling Finger Touch with Fitts' Law. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1363–1372. https://doi.org/10.1145/2470654.2466180

[9] Peter Brandl, Clifton Forlines, Daniel Wigdor, Michael Haller, and Chia Shen. 2008. Combining and Measuring the Benefits of Bimanual Pen and Direct-touch Interaction on Horizontal Interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '08)*. ACM, New York, NY, USA, 154–161. https://doi.org/10.1145/1385569.1385595

[10] Daniel Buschek and Florian Alt. 2015. TouchML: A Machine Learning Toolkit for Modelling Spatial Touch Targeting Behaviour. In *Proceedings of the 20th International Conference on Intelligent User Interfaces (IUI '15)*. ACM, New York, NY, USA, 110–114. https://doi.org/10.1145/2678025.2701381

[11] Daniel Buschek and Florian Alt. 2017. ProbUI: Generalising Touch Target Representations to Enable Declarative Gesture Definition for Probabilistic GUIs. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 4640–4653. https://doi.org/10.1145/3025453.3025502

[12] Daniel Buschek, Alexander Auch, and Florian Alt. 2015. A Toolkit for Analysis and Prediction of Touch Targeting Behaviour on Mobile Websites. In *Proceedings of the 7th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '15)*. ACM, New York, NY, USA, 54–63. https://doi.org/10.1145/2774225.2774851

[13] Daniel Buschek, Alexander De Luca, and Florian Alt. 2015. Improving Accuracy, Applicability and Usability of Keystroke Biometrics on Mobile Touchscreen Devices. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1393–1402. https://doi.org/10.1145/2702123.2702252

[14] Daniel Buschek, Alexander De Luca, and Florian Alt. 2016. Evaluating the Influence of Targets and Hand Postures on Touch-based Behavioural Biometrics. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1349–1361. https://doi.org/10.1145/2858036.2858165

[15] Daniel Buschek, Simon Rogers, and Roderick Murray-Smith. 2013. User-specific Touch Models in a Cross-device Context. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '13)*. ACM, New York, NY, USA, 382–391. https://doi.org/10.1145/2493190.2493206

[16] Matt J. Camilleri, Ajith Malige, Jeffrey Fujimoto, and David M. Rempel. 2013. Touch displays: the effects of palm rejection technology on productivity, comfort, biomechanics and positioning. *Ergonomics* 56, 12 (2013), 1850–62. https://doi.org/10.1080/00140139.2013.847211

[17] Youli Chang, Sehi L'Yi, Kyle Koh, and Jinwook Seo. 2015. Understanding Users' Touch Behavior on Large Mobile Touch-Screens and Assisted Targeting by Tilting Gesture. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1499–1508. https://doi.org/10.1145/2702123.2702425

[18] Karen B. Chen, Anne B. Savage, Amrish O. Chourasia, Douglas A. Wiegmann, and Mary E. Sesto. 2013. Touch Screen Performance by Individuals With and Without Motor Control Disabilities. *Appl Ergon.* 4, 2 (2013), 297–302. https://doi.org/10.1016/j.apergo.2012.08.004

[19] A. Cockburn, D. Ahlström, and C. Gutwin. 2012. Understanding performance in touch selections: Tap, drag and radial pointing drag with finger, stylus and mouse. *International Journal of Human-Computer Studies* 70, 3 (2012), 218 – 233. https://doi.org/10.1016/j.ijhcs.2011.11.002

[20] T. L. Dimond. 1958. Devices for Reading Handwritten Characters. In *Papers and Discussions Presented at the December 9-13, 1957, Eastern Joint Computer Conference: Computers with Deadlines to Meet (IRE-ACM-AIEE '57 (Eastern))*. ACM, New York, NY, USA, 232–237. https://doi.org/10.1145/1457720.1457765

[21] Mark Dunlop and John Levine. 2012. Multidimensional Pareto Optimization of Touchscreen Keyboards for Speed, Familiarity and Improved Spell Checking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2669–2678. https://doi.org/10.1145/2207676.2208659

[22] Mayank Goel, Leah Findlater, and Jacob Wobbrock. 2012. WalkType: Using Accelerometer Data to Accomodate Situational Impairments in Mobile Touch Screen Text Entry. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2687–2696. https://doi.org/10.1145/2207676.2208662

[23] Mayank Goel, Jacob Wobbrock, and Shwetak Patel. 2012. GripSense: Using Built-in Sensors to Detect Hand Posture and Pressure on Commodity Mobile Phones. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, New York, NY, USA, 545–554. https://doi.org/10.1145/2380116.2380184

[24] Joshua Goodman, Gina Venolia, Keith Steury, and Chauncey Parker. 2002. Language Modeling for Soft Keyboards. In *Eighteenth National Conference on Artificial Intelligence*. American Association for Artificial Intelligence, Menlo Park, CA, USA, 419–424. http://dl.acm.org/citation.cfm?id=777092.777159

[25] Asela Gunawardana, Tim Paek, and Christopher Meek. 2010. Usability Guided Key-target Resizing for Soft Keyboards. In *Proceedings of the 15th International Conference on Intelligent User Interfaces (IUI '10)*. ACM, New York, NY, USA, 111–118. https://doi.org/10.1145/1719970.1719986

[26] Niels Henze, Enrico Rukzio, and Susanne Boll. 2011. 100,000,000 Taps: Analysis and Improvement of Touch Performance in the Large. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 133–142. https://doi.org/10.1145/2037373.2037395

[27] Ken Hinckley, Seongkook Heo, Michel Pahud, Christian Holz, Hrvoje Benko, Abigail Sellen, Richard Banks, Kenton O'Hara, Gavin Smyth, and William Buxton. 2016. Pre-Touch Sensing for Mobile Interaction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2869–2881. https://doi.org/10.1145/2858036.2858095

[28] Ken Hinckley, Michel Pahud, Hrvoje Benko, Pourang Irani, François Guimbretière, Marcel Gavriliu, Xiang 'Anthony' Chen, Fabrice Matulic, William Buxton, and Andrew Wilson. 2014. Sensing Techniques for Tablet+Stylus Interaction. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, New York, NY, USA, 605–614. https://doi.org/10.1145/2642918.2647379

[29] Ken Hinckley, Koji Yatani, Michel Pahud, Nicole Coddington, Jenny Rodenhouse, Andy Wilson, Hrvoje Benko, and Bill Buxton. 2010. Pen + Touch = New Tools. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 27–36. https://doi.org/10.1145/1866029.1866036

[30] Christian Holz and Patrick Baudisch. 2010. The Generalized Perceived Input Point Model and How to Double Touch Accuracy by Extracting Fingerprints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 581–590. https://doi.org/10.1145/1753326.1753413

[31] Christian Holz and Patrick Baudisch. 2011. Understanding Touch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2501–2510. https://doi.org/10.1145/1978942.1979308

[32] Christian Holz, Senaka Buthpitiya, and Marius Knaust. 2015. Bodyprint: Biometric User Identification on Mobile Devices Using the Capacitive Touchscreen to Scan Body Parts. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3011–3014. https://doi.org/10.1145/2702123.2702518

[33] Andreas Holzinger, Martin Höller, Martin Schedlbauer, and Berndt Urlesberger. 2008. An investigation of finger versus stylus input in medical scenarios. In *Proceedings of the ITI 2008 30th International Conference on Information Technology Interfaces*. 433–438. https://doi.org/10.1109/ITI.2008.4588449

[34] Sungjae Hwang, Andrea Bianchi, Myungwook Ahn, and Kwangyun Wohn. 2013. MagPen: Magnetically Driven Pen Interactions on and Around Conventional Smartphones. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '13)*. ACM, New York, NY, USA, 412–415. https://doi.org/10.1145/2493190.2493194

[35] Paul Kabbash, I. Scott MacKenzie, and William Buxton. 1993. Human Performance Using Computer Input Devices in the Preferred and Non-preferred Hands. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*. ACM, New York, NY, USA, 474–481. https://doi.org/10.1145/169059.169414

[36] Per-Ola Kristensson and Shumin Zhai. 2004. SHARK2: A Large Vocabulary Shorthand Writing System for Pen-based Computers. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 43–52. https://doi.org/10.1145/1029632.1029640

[37] Per-Ola Kristensson and Shumin Zhai. 2005. Relaxing Stylus Typing Precision by Geometric Pattern Matching. In *Proceedings of the 10th International Conference on Intelligent User Interfaces (IUI '05)*. ACM, New York, NY, USA, 151–158. https://doi.org/10.1145/1040830.1040867

[38] Seungyon Lee and Shumin Zhai. 2009. The Performance of Touch Screen Soft Buttons. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 309–318. https://doi.org/10.1145/1518701.1518750

[39] Markus Löchtefeld, Phillip Schardt, Antonio Krüger, and Sebastian Boring. 2015. Detecting Users Handedness for Ergonomic Adaptation of Mobile User Interfaces. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (MUM '15)*. ACM, New York, NY, USA, 245–249. https://doi.org/10.1145/2836041.2836066

[40] Robert Mack and Kathy Lang. 1989. A Benchmark Comparison of Mouse and Touch Interface Techniques for an Intelligent Workstation Windowing Environment. *Proceedings of the Human Factors Society Annual Meeting* 33, 5 (1989), 325–329. https://doi.org/10.1177/154193128903300520 arXiv:http://dx.doi.org/10.1177/154193128903300520

[41] I. Scott MacKenzie, Abigail Sellen, and William A. S. Buxton. 1991. A Comparison of Input Devices in Element Pointing and Dragging Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '91)*. ACM, New York, NY, USA, 161–166. https://doi.org/10.1145/108844.108868

[42] I. Scott MacKenzie and Shawn X. Zhang. 1999. The Design and Evaluation of a High-performance Soft Keyboard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 25–31. https://doi.org/10.1145/302979.302983

[43] Toshiyuki Masui. 1998. An Efficient Text Input Method for Pen-based Computers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '98)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 328–335. https://doi.org/10.1145/274644.274690

[44] Mohammad Faizuddin Mohd Noor, Andrew Ramsay, Stephen Hughes, Simon Rogers, John Williamson, and Roderick Murray-Smith. 2014. 28 Frames Later: Predicting Screen Touches from Back-of-device Grip Changes. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2005–2008. https://doi.org/10.1145/2556288.2557148

[45] Mohammad Faizuddin Mohd Noor, Simon Rogers, and John Williamson. 2016. Detecting Swipe Errors on Touchscreens Using Grip Modulation. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1909–1920. https://doi.org/10.1145/2858036.2858474

[46] Josip Musić and Roderick Murray-Smith. 2016. Nomadic Input on Mobile Devices: The Influence of Touch Input Technique and Walking Speed on Performance and Offset Modeling. *Human-Computer Interaction* 31, 5 (2016), 420–471. https://doi.org/10.1080/07370024.2015.1071195

[47] Josip Musić, Daryl Weir, Roderick Murray-Smith, and Simon Rogers. 2016. Modelling and correcting for the impact of the gait cycle on touch screen typing accuracy. *mUX: The Journal of Mobile User Experience* 5, 1 (19 Apr 2016), 1. https://doi.org/10.1186/s13678-016-0002-3

[48] Matei Negulescu and Joanna McGrenere. 2015. Grip Change As an Information Side Channel for Mobile Touch Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1519–1522. https://doi.org/10.1145/2702123.2702185

[49] Albert Ng, Michelle Annett, Paul Dietz, Anoop Gupta, and Walter F. Bischof. 2014. In the Blink of an Eye: Investigating Latency Perception During Stylus Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1103–1112. https://doi.org/10.1145/2556288.2557037

[50] Antti Oulasvirta. 2017. User Interface Design with Combinatorial Optimization. *IEEE Computer* 50, 1 (2017), 40–47. https://doi.org/10.1109/MC.2017.6

[51] Antti Oulasvirta, Anna Reichel, Wenbin Li, Yan Zhang, Myroslav Bachynskyi, Keith Vertanen, and Per Ola Kristensson. 2013. Improving Two-thumb Text Entry on Touchscreen Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2765–2774. https://doi.org/10.1145/2470654.2481383

[52] R. L. Potter, L. J. Weldon, and B. Shneiderman. 1988. Improving the Accuracy of Touch Screens: An Experimental Evaluation of Three Strategies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '88)*. ACM, New York, NY, USA, 27–32. https://doi.org/10.1145/57167.57171

[53] Xiangshi Ren and Sachi Mizobuchi. 2005. Investigating the Usability of the Stylus Pen on Handheld Devices. In *SIGHCI 2005 Proceedings*. 12. http://aisel.aisnet.org/sighci2005/12

[54] Xiangshi Ren and Xiaolei Zhou. 2011. An investigation of the usability of the stylus pen for various age groups on personal digital assistants. *Behaviour & IT* 30, 6 (2011), 709–726. https://doi.org/10.1080/01449290903205437

[55] Anne Roudaut, Stéphane Huot, and Eric Lecolinet. 2008. TapTap and MagStick: Improving One-handed Target Acquisition on Small Touch-screens. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '08)*. ACM, New York, NY, USA, 146–153. https://doi.org/10.1145/1385569.1385594

[56] Zhanna Sarsenbayeva, Jorge Goncalves, Juan García, Simon Klakegg, Sirkka Rissanen, Hannu Rintamäki, Jari Hannu, and Vassilis Kostakos. 2016. Situational Impairments to Mobile Interaction in Cold Environments. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 85–96. https://doi.org/10.1145/2971648.2971734

[57] Julia Schwarz, Robert Xiao, Jennifer Mankoff, Scott E. Hudson, and Chris Harrison. 2014. Probabilistic Palm Rejection Using Spatiotemporal Touch Features and Iterative Classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2009–2012. https://doi.org/10.1145/2556288.2557056

[58] Hyunyoung Song, Hrvoje Benko, Francois Guimbretiere, Shahram Izadi, Xiang Cao, and Ken Hinckley. 2011. Grips and Gestures on a Multi-Touch Pen. ACM, 1323–1332. https://www.microsoft.com/en-us/research/publication/grips-and-gestures-on-a-multi-touch-pen/

[59] Ivan E. Sutherland. 1963. *Sketchpad, A Man-Machine Graphical Communication System.* Garland Publishing, New York.

[60] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and Explore Layout Designs with an Optimiser. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16).* ACM, New York, NY, USA, 3780–3783. https://doi.org/10.1145/2851581.2890236

[61] Huawei Tu, Xiangshi Ren, and Shumin Zhai. 2015. Differences and Similarities between Finger and Pen Stroke Gestures on Stationary and Mobile devices. *ACM Trans. Comput.-Hum. Interact.* 22, 5 (2015), 22:1–22:39. https://doi.org/10.1145/2797138

[62] Daniel Vogel and Ravin Balakrishnan. 2010. Occlusion-aware Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10).* ACM, New York, NY, USA, 263–272. https://doi.org/10.1145/1753326.1753365

[63] Daniel Vogel and Patrick Baudisch. 2007. Shift: A Technique for Operating Pen-based Interfaces Using Touch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07).* ACM, New York, NY, USA, 657–666. https://doi.org/10.1145/1240624.1240727

[64] Daniel Vogel, Matthew Cudmore, Géry Casiez, Ravin Balakrishnan, and Liam Keliher. 2009. Hand Occlusion with Tablet-sized Direct Pen Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09).* ACM, New York, NY, USA, 557–566. https://doi.org/10.1145/1518701.1518787

[65] Feng Wang and Xiangshi Ren. 2009. Empirical Evaluation for Finger Input Properties in Multi-touch Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09).* ACM, New York, NY, USA, 1063–1072. https://doi.org/10.1145/1518701.1518864

[66] Daryl Weir, Daniel Buschek, and Simon Rogers. 2013. Sparse Selection of Training Data for Touch Correction Systems. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '13).* ACM, New York, NY, USA, 404–407. https://doi.org/10.1145/2493190.2493241

[67] Daryl Weir, Simon Rogers, Roderick Murray-Smith, and Markus Löchtefeld. 2012. A User-specific Machine Learning Approach for Improving Touch Accuracy on Mobile Devices. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12).* ACM, New York, NY, USA, 465–476. https://doi.org/10.1145/2380116.2380175

[68] Mike Wu, Chia Shen, Kathy Ryall, Clifton Forlines, and Ravin Balakrishnan. 2006. Gesture Registration, Relaxation, and Reuse for Multi-Point Direct-Touch Surfaces. In *First IEEE International Workshop on Horizontal Interactive Human-Computer Systems (Tabletop 2006).* 185–192. https://doi.org/10.1109/TABLETOP.2006.19

[69] Haijun Xia, Tovi Grossman, and George Fitzmaurice. 2015. NanoStylus: Enhancing Input on Ultra-Small Displays with a Finger-Mounted Stylus. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15).* ACM, New York, NY, USA, 447–456. https://doi.org/10.1145/2807442.2807500

[70] Koji Yatani and Khai N. Truong. 2007. An Evaluation of Stylus-based Text Entry Methods on Handheld Devices in Stationary and Mobile Settings. In *Proceedings of the 9th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '07).* ACM, New York, NY, USA, 487–494. https://doi.org/10.1145/1377999.1378059

[71] Ka-Ping Yee. 2004. Two-handed Interaction on a Tablet Display. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04).* ACM, New York, NY, USA, 1493–1496. https://doi.org/10.1145/985921.986098

[72] Ying Yin, Tom Yu Ouyang, Kurt Partridge, and Shumin Zhai. 2013. Making Touchscreen Keyboards Adaptive to Keys, Hand Postures, and Individuals: A Hierarchical Spatial Backoff Model Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13).* ACM, New York, NY, USA, 2775–2784. https://doi.org/10.1145/2470654.2481384

[73] Shumin Zhai and Per-Ola Kristensson. 2003. Shorthand Writing on Stylus Keyboard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03).* ACM, New York, NY, USA, 97–104. https://doi.org/10.1145/642611.642630