

Exploring the Throughput Potential of In-Air Pointing

Michelle A. Brown¹ and Wolfgang Stuerzlinger²

¹ Akendi, Toronto, Canada
michelle@akendi.com

² Simon Fraser University, Burnaby, Canada
<http://ws.iat.sfu.ca>

Abstract. We present an analysis of how pointing performance in in-air un-instrumented pointing can be improved, towards throughput equal to the mouse. Pointing using a chopstick is found to achieve the highest average throughput, with 3.89 bps. This is a substantial improvement over using the finger to point at the screen. Two potential reasons for the throughput gap between chopstick and finger operation were explored: the natural curvature of human fingers and tracking issues that occurs when fingers bend toward the device. Yet, neither one of these factors seems to significantly affect throughput. Thus other, yet unexplored factors must be the cause. Lastly, the effect of unreliable click detection was also explored, as this also affects un-instrumented performance, and was found to have a linear effect.

Keywords: Human-computer interaction · Fitts' law · Pointing tasks · Leap motion

1 Introduction

Un-instrumented in-air interaction has rapidly gained popularity with the introduction of a number of new interaction devices. Potential applications for un-instrumented in-air pointing include interaction in environments where mouse is inadvisable, such as while cooking, mobile computing, medical scenarios [9], and interaction on large wall displays.

The associated tracking devices for in-air interaction enable new and interesting interaction possibilities, including gestures and multi-finger interaction. Yet, previous work [7] has identified that the raw pointing throughput for in-air pointing is substantially less than for the mouse. Thus, it is unclear whether un-instrumented pointing has the potential to match (much less exceed) mouse throughput levels. It is also unclear what aspects of un-instrumented pointing tracking need to be improved to possibly reach mouse-like levels.

Fitts' Law [17] implies that the further away or the smaller a target is, the harder it will be to select. Building on decades of research, the ISO 9241-9 standard [14] standardizes Fitts' law experimental methodologies. It defines throughput T as the primary measure of performance, calculated as $T = \log_2(De/We + 1)/MT$, where, De is the effective distance and We the effective width. These effective values measure the task that the user actually performed, not the one that she or he was presented

with [17]. This reduces variability in identical conditions, which facilitates comparisons between different Fitts' law studies.

1.1 Related Work

Ray pointing is a method for pointing at objects, where the user moves a tracked arm or finger or a tracked object, such as a pen or laser pointer, and orients it in the direction she or he wishes to point to. The first object along that ray is then highlighted and selected when the user indicates selection, e.g., through a button click. Ray pointing remains a popular selection method for large screen and virtual reality systems. Many studies have investigated this technique in large displays [8, 13, 15, 18, 29, 33], Virtual Reality [11, 14, 24, 27], or tabletop scenarios [5]. All these comparisons used devices.

Ray pointing uses 3D input to afford control over a 2D cursor. Effectively users rotate the wrist (or finger) to move the cursor. Early work on finger-pointing used optical tracking [10]. Balakrishnan and MacKenzie [4] identified that a finger affords about 75 % of the bandwidth relative to the wrist. Either moving the finger or the whole hand to control the cursor affords efficient pointing [3]. Yet, tracking very small hand rotations with 3D tracking systems with sufficient accuracy is difficult, as tracking noise is magnified increasingly along the ray. This is the most likely explanation why ray pointing is inferior to other pointing methods in small-scale environments, such as desktops, e.g., [27].

Gallo et al. [9] explored an un-instrumented hand tracking device in a medical context, where sterility is a major concern. Several approaches used various camera systems [14, 19]. In another work [12], the authors look at the requirements of un-instrumented tracking systems and their FingerMouse application used a one-second dwell time for selection. Song et al. [24] used finger pointing to select and move virtual objects. None of the above work evaluates the performance of un-instrumented in-air pointing with the throughput measure as defined by the ISO standard. The exception identified that its throughput was slightly less than 3 bps [7]. This is substantially lower than standard mouse throughput, which is often found to be approximately 4 bps.

1.2 Motivation and Contributions

This paper explores several open explanations for the lower throughput of un-instrumented pointing relative to the mouse [7]. We first evaluated the throughput of a (rigid) chopstick as pointing device, which might have a tracking advantage over a regular finger. It is longer, more cylindrical, and allows for a grip that may offer better directional control. Next, we evaluated pointing throughput of a finger with and without a rigid cast to determine if forcing the finger into a more cylindrical shape would improve tracking. Finally, we investigated the effect of click detection reliability on throughput, as this is another issue that can decrease performance in in-air interaction. Our contributions are:

- An evaluation of the selection performance of a rigid pointing device (chopstick).
- An evaluation of the selection performance of a perfectly cylindrical finger (cast).

- An analysis on the effect unreliable bent finger tracking has on selection performance.
- An evaluation of the effect selection reliability on throughput.

We deliberately chose the Leap Motion for our work, as it is currently one of the best devices for tracking un-instrumented fingers. We considered attaching individual markers to fingers with an optical tracking system. Yet, tracking orientations of fingers requires a large tracking target, which may slow down movements and cause fatigue.

1.3 Pilot Study

Looking at various options to improve tracking robustness, we found that the Leap Motion API supports also long, thin, rigid, cylindrical objects, such as pencils. Based on advice from the Leap Motion forum, we picked a chopstick. We hypothesized that using a chopstick would also increase throughput because it can be held more stably in a pencil grip, i.e., between three fingers.

We recruited 8 participants (mean age 21 years, SD 4.4 years). Two participants were male and all right handed. The Leap Motion sensor was placed directly in front of the display. The Leap Motion software used for this first study was version 1.0.9 + 8410 and the hardware device was LM-010. We used USB3 and Vsync was turned off in all conditions to minimize latency. Both choices increase interaction performance [6], to avoid the potential impact of large differences in latency on pointing performance [22, 24]. End-to-end latency with the Leap Motion was 48 ms, and with the Microsoft IntelliMouse Optical 32 ms. We used the default pointer speed of Windows 7. The software used for this study was FittsStudy [32]. We only added support to read data from the LeapMotion.




Fig. 1. The setup of the pilot study (left), issues observed with tracking bent fingers (right)

For this experiment there were two input conditions for selecting targets for the participants to use: the Chopstick and the Mouse. The Chopstick method required the user to hold a standard disposable wooden chopstick in her or his dominant hand, held like a pencil. Targets were then selected by aiming the tip of the chopstick toward the target on the screen. The Mouse method required the user to operate a computer mouse as they normally would. After targets had been acquired using one of these two methods, targets were selected using the left click button in the Mouse method and the spacebar

on the keyboard in the Chopstick method. The spacebar was operated by the non-dominant hand of the participant and was placed in a comfortable operating position so that the dominant hand used for object acquisition was not obstructed Fig. 1 left illustrates the setup.

First, each participant was given a brief background questionnaire, to record gender, age, and handedness. Then, the participant was introduced to the Chopstick condition and shown how it worked. Participants were required to use a pencil grip for holding the chopstick. After comfortable with basic operation, one of the input conditions was explained to the participant. The order of the input methods was counterbalanced so that each of the possible orders was represented equally. When participants were comfortable with the input method, they completed a series of Fitts' law selection tasks using either the mouse or the chopstick in her or his dominant hand. Ten blocks of 9 Fitts' law conditions with 11 trials were completed with the ISO methodology for a total of 990 trials per condition. Target widths were 32, 64, and 96 pixels and amplitudes 256, 384, and 512 pixels. Then the next input method was presented and the above process repeated. At the end of the experiment, participants were given a brief questionnaire about any discomfort they might have experienced while using un-instrumented tracking and the mouse.

Data was first filtered for obvious participant errors, such as hitting the spacebar twice on the same target or pausing in the middle of a circle (less than .004 % of data collected). For all other analysis and following the ISO standard, we recorded an error whenever the cursor was outside the target upon selection, regardless if this occurred due to human or system, i.e., tracking error. As our data is not normally distributed and fails Levene's test for homogeneity, we conducted ANOVA tests after a Aligned Rank Transform (ART) for nonparametric factorial data analysis, [31].

In terms of throughput there was a significant effect for device used ($F_{1,7} = 19$, $p < .001$) with a power ($1 - \beta$) of .97 and a large effect size (η^2) of .25. For a graph of average throughput values see Fig. 2 (3.54 bps for the chopstick and 4.13 bps for the mouse). There was a significant effect for device used for movement time ($F_{1,7} = 18$, $p < .01$) with a power ($1 - \beta$) of 0.95 and a very small effect size (η^2) of .05. See Fig. 2 for average movement times. There was a significant effect for device used on error rate ($F_{1,7} = 8$, $p < .05$) with a power ($1 - \beta$) of .68 and a negligible effect size (η^2) of .01. The mean error rate was 9.8 % for the chopstick and 3.9 % for the mouse. There was no observed statistically significant learning affect across all blocks ($F_{9,63} = 14$, $p < .001$) with a power ($1 - \beta$) of .99 or in the learning curve between devices ($F_{9,142} = 0.83$, ns). Figure 2 shows performance over time. Device used (chopstick or mouse) crossed with ID value had no significant effect on throughput ($F_{6,97} = 0.02$, ns). Figure 2 shows average movement times for each ID value. The R^2 values show an excellent fit with Fitts' law.

The throughput for the chopstick still has a .39 bps difference in throughput from the mouse by the last block (3.89 vs. 4.28 bps). Yet, latencies in our conditions were in a region (below 50 ms) where they does not seem have a significant effect [22]. This makes it unlikely that latency alone can explain the result. The potential confound of using the mouse and its button with one hand vs. clicking the space bar with the other hand in the chopstick condition is also an unlikely explanation [7]. The error rate for the

chopstick is substantially higher in our current study, either due to limitations in tracking by the Leap Motion or human limits on the ability to point precisely at a distance. Currently we do not have enough information to reliably distinguish between these two causes.




Fig. 2. Graph of average throughput (left) and movement times (middle) for chopstick and mouse. Error bars show standard deviation. Difference is statistically significant. Graph of learning over time (right). Average throughput for each block is displayed. Power curve is fitted to data.

Our results shows that a well-chosen in-air pointing device can achieve high pointing performance: 3.89 bps. That is within the lower end of throughput values observed for the mouse (3.7 bps – 4.9 bps) [25]. With more practice this value may increase further. Interestingly, two participants reached a crossover point where the chopstick achieved a throughput greater than the mouse. An expert user (not a participant), who had been practicing various pointing methods for four months, achieved an average throughput of 4.75 bps with the chopstick and 4.73 bps with the mouse. Yet, while mouse-like levels appear to be attainable with more training, such amounts of training are daunting. Still, we cannot rule out that the chopstick will match the mouse in the long term. In this pilot we did not observe noticeable fatigue effects. The chopstick achieved a throughput of 3.89 bps by the last block, much more than finger operation in prior work [7]. Even accounting for differences in latency (48 ms with our chopstick vs. 63 ms with the finger in [7]), this gap is still substantial. The reason behind this are further explored in the next study.

2 User Study 1

The main objective of this user study was to determine if a perfectly cylindrical, rigid finger would be capable of achieving the same levels of throughput seen with a chopstick in a comparable environment. After all, one possible explanation for the chopstick’s superior performance is its rigid cylindrical nature, making it potentially easier to track. In pilot studies we identified that finger direction tracking reliability of the Leap Motion decreased, if the finger was bent too far towards the tracking device. See Fig. 1, right for a depiction of this problem. In this figure, the top two frames show a straight finger and the corresponding finger direction arrow. Subsequent frames show results with

increasing finger bend, where the direction deviates more and more. Moreover, we observed that some users had significantly more curved fingers than others. An example for this finger curve is visible in the index finger in Fig. 3, rightmost image.



Fig. 3. Pictures of the four input conditions. From top left to bottom right: *Cast Normal*, *Cast Side*, *Normal*, and *Side*

We also speculated that finger tracking might behave differently depending on whether the users held their hands palm facing down or rotated 90° inwards. We included such conditions here as it might be easier for the device to track the position of the finger and determine the pointing direction – if finger curvature plays a significant role.

2.1 Input Conditions

For this user study there were four input conditions for selecting targets for the participants to use. These were the *Cast Normal*, the *Cast Side*, the *Normal*, and the *Side* method, as shown in Fig. 3. The *Cast Normal* method required the user to wear a paper “cast” around her or his dominant pointer finger. This cast was specially designed and adapted to each user’s finger. A piece of regular computer paper was cut so that it was wide enough to wrap around the user’s finger and long enough to cover the finger to the tip. This piece of paper was then wrapped around the user’s finger and taped with clear adhesive tape to form the “cast”. The finger was held in the “normal” pointing orientation with the bottom of the user’s palm facing down. In the *Cast Side* method, the “cast” was again worn on the user’s finger but this time the finger was held in the “side” position with the user’s palm perpendicular to the desk. The *Normal* method required the user to hold their hand with the palm facing down, toward the desk, without a cast. In the *Side* condition the user’s palm was held perpendicular to the desk, again without a cast. In all conditions, after targets had been acquired through pointing, selection was indicated via the spacebar on the keyboard. The spacebar was operated by the non-dominant hand of the participant and was placed in a comfortable operating position so that the dominant hand used for object acquisition was not obstructed. We hypothesized here that if finger cast performance reaches chopstick levels, then the grip style is likely not the cause of the chopstick’s performance. In this case, rigidity would be a more likely explanation.

2.2 Participants and Procedure

We recruited 8 different participants for this study (mean age 20 years, SD 2.3 years). Three participants were male and all but one were right handed. First, participants were given a brief background questionnaire which recorded gender, age, and handedness.

Next, a “cast” was created for each participant as described in the Input Conditions. Then, the participant was introduced to the finger tracking system and the experimenter demonstrated how it worked. After was comfortable with basic operation, one of the input conditions was explained to the participant. The order that participants were exposed to each of the input methods was determined by a Latin Square design. Once comfortable with the current input method, the participant completed a series of Fitts’ law selection tasks using one of the four input conditions. Five blocks of 9 Fitts’ law conditions with 11 trials per condition for a total of 495 trials were completed, again using the ISO methodology. Target widths of 32, 64, and 96 and target amplitudes of 256, 384, and 512 pixels were used. The participant was then presented with the next input method and so on.

2.3 Results

Data was first filtered for errors, such as hitting the spacebar twice on the same target or unusually long pauses (less than .01 % of total data). The data is not normally distributed and fails Levene’s test for homogeneity, and we again used ART before ANOVA.

There was no significant effect for the used interaction method ($F_{3,21} = 1.35, p : .05$) on throughput, nor for any pair of conditions. See Fig. 4 for average throughput values.




Fig. 4. Graph of average throughput values (left) and movement time (middle) and Fitts’ law model (right) for each condition. Error bars show standard deviation.

There was no significance effect for the used interaction method ($F_{3,21} = 2.57, p > .05$) on movement time, nor for any pair of conditions. See Fig. 4 for average movement times. The used interaction method had no significant effect on error rate ($F_{3,21} = 0.27, ns$). The four conditions had error rates of 14 %, 12 %, 13 % and 13 % respectively. Across all blocks there was no significant effect on learning ($F_{4,28} = 1.15, p > .05$) and no effect on learning crossed with the used interaction method ($F_{12,145} = 1.64, p > .05$). The used interaction method crossed with ID had no significant effect on throughput ($F_{18,207} = 1.03, p > .05$). See Fig. 4 for the data for all conditions. The equations and fit values for the Fitts’ law models are as follows: *Cast Normal*: $y = 310.58 x - 7.9434, R^2 = 0.9857$, *Cast Side*: $y = 356.31 x - 26.25$,

$R^2 = 0.9743$, *Normal*: $y = 337x - 49.201$, $R^2 = 0.9826$, *Side*: $y = 329.88x - 8.9465$, $R^2 = 0.9879$, again all conforming to Fitts' law.

2.4 Discussion

This study indicates that the cast conditions are similar to finger tracking. Therefore, it is unlikely that the natural curvedness and potential flexibility of a human finger cause lower pointing throughput relative to a rigid object. Yet, there is still a 15 + % difference (0.6 bps) between the throughputs of finger operation and chopstick operation that remains unaccounted for. The higher throughput from the pilot study must thus be due to some other factor, such as tracking a longer object or the different grip on the chopstick. Our results largely confirm the results of previous work [4], but also extend it through our use of the ISO methodology, which removes the effect of the speed-accuracy tradeoff.

Moreover, informal observations during this experiment identify fatigue as a potential issue, similar to [7]. This may be due to the duration of the experiment, which lasted about one hour. After all, many people are not used to using their index finger for long periods as a pointing “instrument”. Still, performance did not drop noticeably in later trials.

3 User Study 2

To further investigate the potential of in-air interaction, we decided to look at the effect that varying degrees of click detection reliability have on throughput. After all, even a device that affords highly precise pointing may suffer if the selection of targets cannot be indicated reliably. To accurately and *reproducibly* control the level of reliability, we decided to perform this study with a mouse, as its buttons are normally 100 % reliable. The results of such an experiment can then be used to infer the *potential* performance impact of a selection method that is not 100 % reliable, such as in in-air “click”.

3.1 Participants, Setup and Procedure

We recruited 10 different participants for this study (mean age 23 years, SD 4.7 years). Four participants were male and all but one were right handed. The left-handed person preferred to operate the mouse with the right hand. The mouse used was a Microsoft IntelliMouse Optical set to the default pointer speed on the Windows 7 operating system. The system used with the mouse had an end-to-end latency of 28 ms (Vsync was off). The software used for conducting the Fitts' law tasks was again FittsStudy [32].

First, the participant was given a brief background questionnaire to record gender, age, and handedness. Then, the participant was informed that the mouse button used for clicking would not always be reliable and that sometimes it might need to be clicked again. We chose to inform participants in advance to avoid potential confounds due to side effects of frustration. We tested five levels of reliability: 100 %, 99 %, 98 %, 95 %, and 90 %, to keep frustration levels at an acceptable level. The order that participants

received each of these conditions was counterbalanced so that each of the possible orders was represented equally. Participants then completed 2 blocks of 12 Fitts' law conditions with 11 trials per condition for a total of 264 targets with the ISO methodology. Target widths of 16, 32, 64, and 96 pixels and amplitudes of 256, 384, and 512 pixels were used.

3.2 Results

As our data is not normally distributed and fails Levene's test for homogeneity, all ANOVA tests were again conducted on data transformed using ART.

There was a significant effect for reliability level ($F_{4,36} = 7, p < .001$) on throughput with a power ($1 - \beta$) of .99 and a medium effect size (η^2) of .09. A Tukey-Kramer Multiple-Comparison test identified two statistically different groups. Group one consists of 90 % and 95 % reliability and group two of 98 %, 99 %, and 100 % reliability. See Fig. 5 for average throughput values. There was a significant effect for reliability level ($F_{4,36} = 8, p < .001$) on movement time with a power ($1 - \beta$) of .99 and a very small effect size (η^2) of .01. A Tukey-Kramer Multiple-Comparison test again identified two statistically significant groupings. However, the groupings were different than the throughput groupings. Group one consisted of 90 %, 95 % and 98 % reliability and group two consisted of 98 %, 99 % and 100 % reliability. In other words, 98 % was not statistically different from all other conditions. See Fig. 5 for average movement times.

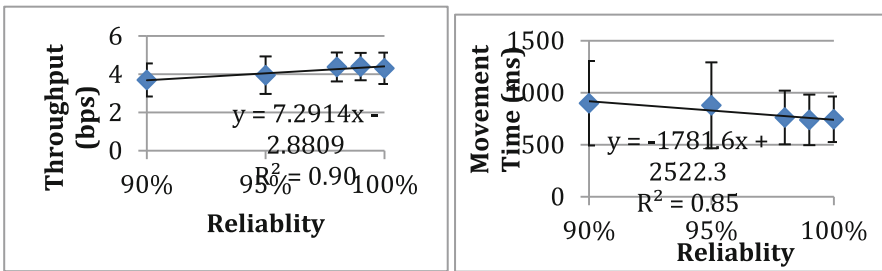


Fig. 5. Average throughput values (left) and movement times (right) for each reliability level. Error bars show standard deviation. A linear trendline and its corresponding equation is also shown.

The used reliability level had no significant effect on error rate ($F_{4,9} = 1.86, p > .05$). The mean error rates for the 90 % to 100 % conditions were 4.2 %, 1.4 %, 3.7 %, 4.8 % and 0.8 % respectively. Across all blocks, there was no significant effect on learning ($F_{1,18} = 0.05, ns$) and no effect on learning crossed with level of reliability ($F_{4,85} = 1.08, p > .05$). Reliability level crossed with ID had no significant effect on throughput ($F_{32,428} = 6, p > .05$). See Fig. 5 for the data for all conditions. The equations for the Fitts' law models are as follows: 90 %: $y = 170.82x + 155.98, R^2 = 0.988$, 95 %: $y = 148.78x + 212.89, R^2 = 0.988$, 98 %: $y = 167.44x + 95.258, R^2 = 0.987$, 99 %: $y = 156.39x + 103.94, R^2 = 0.996$, 100 %: $y = 141.48x + 138.47, R^2 = 0.998$.

3.3 Discussion

These results indicate that there is a roughly linear drop-off in pointing performance as a selection technique becomes more unreliable. The 90 % and 95 % conditions performed significantly worse than 98 % and above in terms of throughput. We see this as an indication (but not as proof) that any click-gesture recognition system that is 95 % reliable or less is going to noticeably and negatively impact interaction performance with a system. While there was no significant difference in performance between 100 %, 99 %, and 98 %, some participants did still notice when they were not at 100 % condition. This indicates that while a system with reliability above 95 % might not suffer much in terms of throughput, failures might still be noticeable to the users. Small amounts of errors might be less notable in systems without force feedback or where users expect it to be unreliable.

From observations during the experiment we also identified a behavioural difference for many in the 90 % condition: most participants would pause after selecting a target before the next one. Thus, it seemed like the participants expected failure rather than success in the 90 % condition. We suspect that as the reliability gets even lower all participants would anticipate a failure, not just most of them.

Perfect reliably in un-instrumented in-air pointing with a single camera is very difficult. Even very recent work does not achieve 100 % reliability [23]. Thus on top of tracking issues one must also factor in a loss in throughput due to click detection unreliability.

4 Overall Discussion

We explored several possibilities for the lower throughput of un-instrumented pointing relative to the mouse, as identified by previous work [7]. First, we identified that pointing with chopsticks can approach the performance traditionally seen with mice. This points to new interesting avenues for future user interfaces. We also evaluated finger pointing with and without a rigid cast. Given that we found no significant difference, it is unlikely that the rigidity of the input device is the primary explanation. This leaves the length of the chopstick or the grip style as possible explanations. Finally, we evaluated the effect of click detection reliability on throughput, another potential issue in in-air interaction. Our results indicate that in-air “click” detection must have between 95 and 98 % reliability, for in-air interaction to have the potential to perform as well as a mouse.

5 Conclusion

In this paper we evaluated several factors that were hypothesized to affect pointing performance: the shape of the finger, finger bend tracking difficulties, and click detection reliability. Moreover, we showed that by using a chopstick, users could reach the lower end of the range of pointing throughputs seen with the mouse. We also identified that finger curvedness or rigidity have no effect on pointing throughput with the Leap Motion.

Finally, we showed that unreliable selection techniques affect performance (approximately) linearly and identified key values between 90 % and 100 % reliability.

References

1. Arif, A.S., Stuerzlinger, W.: Predicting the cost of error correction in character-based text entry technologies. In: ACM CHI 2010, pp. 5–14 (2010)
2. Arif, A.S., Stuerzlinger, W.: User adaptation to a faulty unistroke-based text entry technique by switching to an alternative gesture set. *Graphics Interface 2014*, pp. 183–192 (2014)
3. Balakrishnan, R., Baude, T., Kurtenbach, G., Fitzmaurice, G.: The Rockin’Mouse: integral 3D manipulation on a plane. In: ACM CHI 1997, pp. 311–318 (1997)
4. Balakrishnan, R., MacKenzie, I.S.: Performance differences in the fingers, wrist, and forearm in computer input control. In: ACM CHI 1997, pp. 303–310 (1997)
5. Banerjee, A., Burstyn, J.: Pointable: an in-air pointing technique to manipulate out-of-reach targets on tabletops. In: ACM ITS 2011, pp. 11–20 (2011)
6. Bedikian, R.: “Understanding Latency,” Leap Motion Developer Labs (2013). <http://labs.leapmotion.com/post/55354675113/understanding-latency-part-1>
7. Brown, M.A., Stuerzlinger, W., Mendonça Filho, E.J.: The performance of un-instrumented in-air pointing. In: *Graphics Interface 2014*, pp. 59–66 (2014)
8. Das, K., Borst, C.W.: An evaluation of menu properties and pointing techniques in a projection-based VR environment. In: IEEE 3DUI 2010, pp. 47–50 (2010)
9. Gallo, L., Placitelli, A., Ciampi, M.: Controller-free exploration of medical image data: experiencing the kinect. In: *Computer-Based Medical Systems 2011*, pp. 1–6 (2011)
10. Gokturk, M., Sibert, J.L.: An analysis of the index finger as a pointing device. In: *Extended Abstracts ACM CHI 1999*, p. 286 (1999)
11. Grossman, T., Balakrishnan, R.: Pointing at trivariate targets in 3D environments. In: ACM CHI 2004, pp. 447–454 (2004)
12. Von Hardenberg, C., Bérard, F.: Bare-hand human-computer interaction. In: *Workshop on Perceptive User Interfaces 2001* (2001)
13. Jota, R., Nacenta, M., Jorge, J.: A comparison of ray pointing techniques for very large displays. In: *Graphics Interface 2010*, pp. 269–276 (2010)
14. Kolaric, S., Raposo, A., Gattass, M.: Direct 3D manipulation using vision-based recognition of uninstrumented hands. In: *Symposium on Virtual and Augmented Reality 2008*, pp. 212–220 (2008)
15. Kunert, A., Kulik, A., Lux, C., Fröhlich, B.: Facilitating system control in ray-based interaction tasks. In: ACM Symposium VRST 2009, pp. 183–186 (2009)
16. ISO 9241-9 Ergonomic requirements for office work with visual display terminals (VDTs)-Part 9: Requirements for non-keyboard input devices, ISO (2000)
17. MacKenzie, I.S.: Fitts’ law as a research and design tool in human-computer interaction. *Hum. Comput. Interact.* 7(1), 91–139 (1992)
18. MacKenzie, I., Jusoh, S.: An evaluation of two input devices for remote pointing. In: Nigay, L., Little, M. (eds.) *EHCI 2001*. LNCS, vol. 2254, pp. 235–250. Springer, Heidelberg (2001)
19. Matikainen, P., Pillai, P., Mummert, L., Sukthankar, R., Hebert, M.: Prop-free pointing detection in dynamic cluttered environments. In: *Face and Gesture 2011*, pp. 374–381 (2011)
20. Oakley, I., Sunwoo, J., Cho, I.-Y.: Pointing with fingers, hands and arms for wearable computing. In: *Extended Abstracts ACM CHI 2008*, pp. 3255–3260 (2008)
21. Oh, J., Stuerzlinger, W.: Laser pointers as collaborative pointing devices. In: *Graphics Interface 2002*, pp. 141–149 (2002)

22. Pavlovych, A., Stuerzlinger, W.: The tradeoff between spatial jitter and latency in pointing tasks. In: ACM EICS 2009, pp. 187–196 (2009)
23. Sharp, T., Keskin, C., Robertson, D., Taylor, J., Shotton, J., Kim, D., Rhemann, C., Leichter, I., Vinnikov, A., Wei, Y., Freedman, D., Kohli, P., Krupka, E., Fitzgibbon, A., and Izadi, S.: Accurate, robust, and flexible real-time hand tracking. In: ACM CHI 2015, pp. 3633–3642 (2015)
24. Song, P., Yu, H., Winkler, S.: Vision-based 3D finger interactions for mixed reality games with physics simulation. In: ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry 2008, article no. 7 (2008)
25. Soukoreff, R.W., MacKenzie, I.S.: Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *IJHCS* **61**(6), 751–789 (2004)
26. Teather, R.J., Pavlovych, A., Stuerzlinger, W. MacKenzie, I.S.: Effects of tracking technology, latency, and spatial jitter on object movement. In: IEEE Symposium 3DUI 2009, pp. 43–50 (2009)
27. Teather, R.J., Stuerzlinger, W.: Pointing at 3D targets in a stereo head-tracked virtual environment. In: IEEE Symposium 3DUI 2011, pp. 87–94 (2011)
28. Teather, R.J., Stuerzlinger, W.: Pointing at 3D target projections with one-eyed and stereo cursors. In: ACM CHI 2013, pp. 159–168 (2013)
29. Vogel, D., Balakrishnan, R.: Distant freehand pointing and clicking on very large, high resolution displays. In: ACM UIST 2005, pp. 33–42 (2005)
30. Wingrave, C., Bowman, D.: Baseline factors for raycasting selection. In: HCI International 2005 (2005)
31. Wobbrock, J.O., Findlater, L., Gergle, D., Higgins, J.J.: The aligned rank transform for nonparametric factorial analyses using only ANOVA procedures. In: ACM CHI 2011, pp. 143–146 (2011)
32. Wobbrock, J.O., Shinohara, K., Jansen, A.: The effects of task dimensionality, endpoint deviation, throughput calculation, and experiment design on pointing measures and models. In: ACM CHI 2011, pp. 1639–1648 (2011)
33. Zigelbaum, J., Browning, A., Leithinger, D., Bau, O., Ishii, H.: G-stalt: a chirocentric, spatiotemporal, and telekinetic gestural interface. In: ACM TEI 2010, pp. 261–264 (2010)