

How low can you go? Human limits in small unidirectional mouse movements

Jonathan Aceituno, Géry Casiez, Nicolas Roussel

► **To cite this version:**

Jonathan Aceituno, Géry Casiez, Nicolas Roussel. How low can you go? Human limits in small unidirectional mouse movements. CHI'13, the 31th Conference on Human Factors in Computing Systems, Apr 2013, Paris, France. pp.1383-1386, 2013, <10.1145/2470654.2466182>. <hal-00799954>

HAL Id: hal-00799954

<https://hal.inria.fr/hal-00799954>

Submitted on 3 May 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How Low Can You Go?

Human Limits in Small Unidirectional Mouse Movements

Jonathan Aceituno¹, Géry Casiez^{1,2,3} & Nicolas Roussel¹

¹Inria Lille & ²LIFL, ³University of Lille, France
jonathan.aceituno@inria.fr, gery.casiez@lfl.fr, nicolas.roussel@inria.fr

ABSTRACT

Computer mouse sensors keep increasing in resolution. The smallest displacement they can detect gets smaller, but little is known on our ability to control such small movements. Small target acquisition has been previously tackled, but the findings do not apply to the problem of finding the *useful resolution* of a user with a mouse, which corresponds to the smallest displacement (s)he can reliably produce with that device. We detail this definition and provide an associated experimental protocol to measure it. We then report on the results of a study suggesting that high-end mice are not likely to be used to their full potential. We further comment on the different strategies used by participants to achieve best performance, and derive implications for user interfaces.

ACM Classification Keywords

H.5.2 [Information interfaces and presentation]: User interfaces - Input devices and strategies.

General Terms

Experimentation, Human Factors

Author Keywords

Input device; mouse; resolution; useful resolution

INTRODUCTION

The resolution of a sensor is "*the smallest detectable change of input signal that causes a change in output signal*" [7]. For a computer mouse, it is the smallest measurable displacement. The length of this displacement is rarely reported, however. Instead, it is usually referred to as one *count* and the resolution given as the number of counts per inch (CPI). The resolution of modern mice typically ranges from 400 to 10000 CPI, which corresponds to displacements between $64\mu\text{m}$ and $2.54\mu\text{m}$ (about the size of a bacterium). Today's high-end mice are thus probably among the cheapest and most prevalent high-resolution sensors in the world.

Mouse manufacturers sometimes use high-resolution as a selling argument, especially for gaming products. But, to our knowledge, human ability to control small displacements with these high-end devices has never been studied. Probably due to the lack of high-resolution sensors until recently,

little is actually known about the motor resolution of the human hand in indirect pointing situations. Indirect high-precision tasks can of course be eased by means of transfer functions using low control-display gain (*CD gain*) values to decrease in display space the amplitude of movements in motor space [2]. But can users actually take advantage of a 10000 CPI mouse? Is the quest for higher resolutions worth pursuing? Are the small pointer movements required by today's systems achievable without assistance?

Although there is no specific literature on human ability to control small displacements of high-end mice, related information can be found in Fitts' law studies involving small target widths in motor space. Langolf et al. reported successful acquisitions of 0.076 mm targets in a peg transfer task under a stereoscopic microscope, for example [5]. In a multi-scale pointing task, Guiard et al. reported successful acquisitions of 0.06 mm targets using a Wacom digitizer with a stylus and a puck [4]. Chapuis and Dragicevic also reported successful acquisitions of 0.06 mm targets using an "ultra-high-resolution gaming mouse of 83.5 dots per mm (about 2000 DPI)" [3]. Bérard et al. more recently found that participants could comfortably acquire 0.036 mm and 0.018 mm targets with a mouse, and even smaller ones (0.0045 mm) without impeding the error rate but at the expense of longer movement times [1]. But acquiring a small target at a certain distance is not necessarily the same as generating an equivalently small displacement. It is not clear whether the above resolutions remain valid for tasks other than target acquisition, like finely adjusting an object's position for example.

Our focus in this work is on the *useful resolution* of a user equipped with a particular device. Our interest is not in the smallest target (s)he can acquire, but in *the smallest displacement (s)he can reliably produce*. After clarifying the important differences between the two approaches, we propose an experimental protocol to measure the useful resolution based on series of increasingly small controlled movements. We then report on the results of an experiment based on this protocol and discuss their implications.

DETERMINING THE USEFUL RESOLUTION

The aim of this work is to complement current knowledge on the issues related to small-scale control of input devices. We identify two main aspects of movement where human limits can apply: target sizes and movement amplitude. Although the former aspect has been previously investigated, we demonstrate that the findings do not apply to the problem of determining a lower limit regarding movement amplitudes.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI'13, April 27–May 2, 2013, Paris, France.

Copyright 2013 ACM 978-1-4503-1899-0/13/04...\$15.00.

We then define that limit as the *useful resolution* and discuss its measurement.

Pointing is not the point

To our knowledge, the study reported by Bérard et al. [1] is the only one that explicitly investigated the lower limit of target sizes in motor space. More specifically, they wanted to determine "the smallest target size that users can acquire with an ordinary amount of effort using a particular device". Their study shows that this size depends on the form factor of the device, and more specifically on its self-stabilizing property.

Pointing tasks are used to study small-scale target acquisition with mice, but are they appropriate for determining the lower limits of controlled movements? Meyer et al's optimized dual-submovement model describes pointing as involving a primary ballistic phase followed by an optional corrective one [6]. As target size decreases, the ballistic submovement is more likely to fall on either side of the target and thus require corrective submovements (Figure 1). It is however unclear whether these ballistic and corrective submovements bear any relation to the smallest controllable movement. If one wants to characterize a small movement, one needs to control the conditions under which it can be produced.




Figure 1: Ballistic and corrective phases in pointing tasks.

From device resolution to useful resolution

For a pointing device of resolution R , the smallest reportable displacement is $d = 1/R$ (in inches) and any displacement in $[n \times d; (n + 1) \times d]$ results in a report of n counts, assuming that the device sensors did not accumulate any previous displacement. The *user resolution* corresponding to a displacement of p counts reported by this device is thus $R_p = 1/(p+1) \times d = R/(p+1)$ (in CPI). From this equation, it is clear that the user resolution space for this device is a discrete one with a maximum value of $R_1 = R/2$ (Table 1).

$p_{(\text{counts})}$	1	2	3	4	5	6	7	15	31	63
$R_p_{(\text{CPI})}$	3200	2133	1600	1280	1066	914	800	400	200	100

Table 1: Partial view of the user resolution space of a 6400 CPI mouse.

As explained in the introduction, the *useful resolution* of a user with a device is the smallest displacement (s)he can reliably produce with it. Determining this resolution thus consists in finding the lowest p value (s)he can reliably produce. To determine this value, we propose an experimental protocol following a *limbo* approach. The protocol works as follows:

1. Choose a starting p value (a number of counts ≥ 1).
2. Test R_p : ask the user to move the device so that it reports at most p counts. Do this k times and compute the success

rate S_p defined as the percentage of repetitions where the produced counts were at most p .

A user will be considered able to reliably produce p counts if S_p exceeds a threshold α . (S)he will be considered unable to produce these counts if S_p falls below a second threshold β .

3. If $p > 1$ and $S_p > \beta$, choose a lower p value and go back to step 2.
4. The useful resolution \hat{R} is the greatest R_p for which $S_p \geq \alpha$.

As p decreases, S_p is expected to monotonically decrease. The starting p value should be chosen so as to yield high success rates with a high probability. The choice of subsequent p values can be guided by the user resolution space of the device (e.g. Table 1 for a 6400 CPI mouse). The α threshold corresponds to the expected degree of reliability. The β threshold determines the stopping condition of the protocol (e.g. run until $p = 1$ for $\beta = 0$, or stop after finding the first \hat{R} candidate for $\beta = \alpha$).

EXPERIMENT

We conducted an experiment to study the useful resolution for a particular mouse and the potential influence of movement direction on it. To this end, we followed the above protocol using $\beta = 0$, and deliberately not choosing α .

Apparatus, task and participants

We used a right-handed Razer Imperator 2012 mouse on a varnished plywood desk (Figure 2a, coefficient of static friction $\mu = 0.17$). We used libpointing [2] to bypass the system's transfer function and access raw motion information. We adjusted the polling rate of the mouse to 500 Hz and verified its advertised resolution (6400 CPI) by looking at the raw displacement reported for a reference distance measured with a ruler. We also used Razer software to calibrate the mouse for the desk's surface. The display was a 22" LCD monitor (DELL 2208WFP) using a 1680×1050 pixel resolution (90 PPI) at 60 Hz. Programmed in C++ and JavaScript using WebKit, the experiment ran full-screen on an Intel Core i7 MacBook Air.

Based on the above protocol, the task consisted in moving the mouse along a requested direction for a distance of at least 1 count and up to a specified maximum value. To present all conditions in a consistent way and avoid visual acuity related problems, the maximum value and the traveled distance were shown as numbers, in counts (Figure 2). The requested direction was presented as an arrow. Participants were instructed to do their best to stay under the maximum allowed distance, without any time or posture constraint.

Trials started with the first displacement report received from the mouse and ended either because the accumulated traveled distance exceeded the maximum allowed one (failed trial), or after 750 ms without receiving any report (successful one). Determined through pilot tests, this timeout for movement segmentation was long enough to avoid false positives caused by uncontrolled pauses and short enough to avoid false negatives caused by uncontrolled movements after task completion. Any motion reported in the opposite direction to the requested one canceled the ongoing trial which automatically restarted after a 750 ms pause explaining the situation. Upon




Figure 2: Physical layout (left) and display for failed (top-right) and successful (bottom-right) trials.

completion, a message indicating whether the trial was successful or not was displayed for 750 ms, after which the instructions for the next trial were automatically presented.

Twelve unpaid volunteers with a mean age of 24.6 ($SD = 2.7$) served in the experiment (8 male, 4 female, all right-handed). Seven used a computer more than 6 hours a day. Three were using a mouse in computer games more than 2 hours a day. None suffered from any visuo-motor impairment.

Design

A repeated-measures within-subjects design was used. The independent variables were the movement direction (DIR) and the maximum allowed movement distance in that direction, expressed as a resolution (RES) to facilitate further discussion. DIR was evaluated with four levels (EAST, WEST, NORTH and SOUTH) aligned with the mouse axes. Eight levels of RES were presented: 100, 200, 400, 800, 1280, 1600, 2133 and 3200 CPI.

Participants were given a few minutes to get used to the task before starting the experiment. Then they completed three successive BLOCKS. Each BLOCK consisted of 192 trials: 6 repetitions of the 32 DIR \times RES combinations. RES were presented in ascending order. The presentation order for DIR was counterbalanced across participants using a Latin Square design. Participants had to press a key after each series of 6 trials to move to the next, and were encouraged to take a break before doing so. The experiment lasted approximately 40 minutes. At the end of it, participants were interviewed.

In summary, the experimental design was: 12 participants \times 3 BLOCKS \times 4 DIRECTIONS \times 8 RESOLUTIONS \times 6 trials = 6912 total trials.

RESULTS

The dependent variable is the success rate. Canceled trials (10.5%) were filtered out for the analysis. The first trial for each block was also removed as we observed participants sometimes did not notice the condition changed. As success rate exhibited a non-normal distribution, data were pre-processed using an Aligned Rank Transform [9]. We then ran a repeated-measures ANOVA and tested for significant interactions between factors¹.

¹We used Bonferroni correction for pairwise comparisons, and when the assumption of sphericity was violated, we corrected the degrees of freedom using Greenhouse-Geisser estimates of sphericity.




Figure 3: Mean success rate for RES and DIR. Error bars represent 95% confidence interval.

We found a significant effect of BLOCK on success rate ($F_{2,22}=18.6$, $p<0.001$). Pairwise comparisons showed a significant increase of success rate between the first block and the two remaining ($p<0.05$; Block 1: 54.3%, Block 2: 64.6%, Block 3: 67.1%), indicating a learning effect. We thus removed the first block from subsequent analyses.

DIR was found to have a significant main effect on success rate ($F_{3,33}=11.5$, $p<0.001$, Figure 3). Pairwise comparisons showed that NORTH was significantly different from EAST ($p<0.05$) and WEST ($p<0.05$), and that there was a marginally significant difference between SOUTH and EAST ($p=0.055$), but no other significant difference was found (EAST: 74.4%, NORTH: 52.7%, SOUTH: 62.5%, WEST: 73.5%).

We also found a significant main effect of RES ($F_{3,5,38,3}=102.2$, $p<0.001$) and significant DIR \times RES interaction ($F_{6,3,68,9}=3.4$, $p<0.05$) on success rate. Pairwise comparisons did not show significant differences in success rate between directions for 100 and 200 CPI. However, for 400 CPI, success rate was found significantly lower for NORTH than for WEST ($p=0.037$). For 800 CPI, we found that success rate was significantly lower for NORTH than for EAST ($p=0.046$) and for WEST ($p=0.010$). For 1280 CPI, we found success rate to be significantly lower for NORTH than for EAST ($p=0.008$) and for WEST ($p=0.031$). For 1600 CPI, success rate was significantly higher for EAST than for SOUTH ($p=0.030$), and it was significantly lower for NORTH than for WEST ($p=0.004$). For 2133 CPI and 3200 CPI, no significant difference between directions was found on success rate.

To compare success rates between resolutions, we removed for each one the directions with significantly lower success rate and aggregated the remaining directions. Repeated measures ANOVA showed a significant main effect of RES on success rate ($F_{2,8,30,9}=31.4$, $p<0.001$). Pairwise comparisons showed that 100 and 200 CPI were significantly different from the other resolutions ($p<0.03$). The average success rate for resolutions 100 and 200 CPI is 95.5%. Note that we found the same significant effects when keeping all directions.

DISCUSSION

Our analysis did not show any difference across directions for 100 and 200 CPI, for which participants were able to successfully complete the task 95.5% of the time. Significant differences were found, however, between the CPI ranges [100-200] and [400-3200]. Our results suggest that the useful resolution with the tested device falls within the range [200-400] CPI for most participants. Using $\alpha = 95\%$, we computed in-

dividual values. \hat{R} was 100 CPI for 3 participants, 200 for 4, 400 for 4 others and 1280 for the last one. Recalling that the maximum user resolution is half that of the device, this means a 400 CPI mouse similar in every other characteristics to the tested one would be sufficient for 7 of our 12 participants. An 800 CPI one would be sufficient for 11 of them, and even our most precise participant would not benefit from a 3200 CPI one. Though other reasons may prevail, the quest for mouse resolutions above 10000 CPI does not seem worth pursuing from a useful resolution perspective.

Participant strategies

Participants were asked to describe the main strategy they employed to complete the tasks. Three participants simply used the device as usual (2 with $\hat{R} = 100$ CPI, 1 with 400 CPI). The most used strategy consisted in blocking the mouse using a part of the hand while pushing it along the requested direction and slowly releasing the blocking to achieve fine control (seven participants, 4 with $\hat{R} = 200$ CPI, 3 with 400 CPI). When going WEST, for example, they would block the device using the thumb. Instead of moving in straight line, one participant ($\hat{R} = 100$ CPI) performed small rotations whose center were seldom aligned with the device sensor, resulting in more canceled trials. One participant ($\hat{R} = 1280$ CPI) performed diagonal movements.

The fact that most participants used the block-and-push strategy could explain the significantly lower success rate observed for NORTH since no part of the hand can be used to block the device in that direction (Figure 2). For the other directions, we hypothesize block-and-push helps overcoming stick-slip phenomena when the exerted force exceeds static friction between the mouse and the desk. Should this hypothesis be true, software alternatives could be implemented such as not taking into account the first counts reported for a movement. Further experiments are required to fully characterize the different strategies and the influence of static friction.

Implications for user interfaces

The range of useful resolutions we found overlaps that of modern displays, typically between 100 to 200 PPI. As the default transfer functions of modern systems produce CD gain values close to 1 at low speeds [2], pixel-precise tasks should not require further assistance. However, in situations where the useful resolution is higher than the display resolution, one could imagine taking advantage of the higher precision for tasks where pixels are not enough, provided that additional feedback is given. For instance, using a conventional slider would allow frame-by-frame navigation in a video [8]. A tool for estimating one's useful resolution with a particular device could provide systems and applications with the necessary information. Such a tool could follow the proposed protocol with a starting p value corresponding to 200 CPI, which seems a reasonable middle ground for most users, and an α value of 95%. For most users, based on our results, iterating two or three times through the protocol should be enough.

CONCLUSION

We introduced the *useful resolution*, which we defined as the smallest displacement a user can reliably produce with

a particular device. After explaining how this concept departs from the existing literature on small target acquisition, we proposed an experimental protocol to measure it based on series of increasingly small controlled movements. This concept and the associated protocol are relevant to any pointing device. We reported on an experiment based on this protocol for a high-end mouse. Our results suggest choosing 95% as the threshold for the reliable production of small displacements and show that most participants had a useful resolution between 200 and 400 CPI with the tested device. Participants reported using different strategies to achieve best performance, and we conjecture the most used one played a role in overcoming the effect of static friction. We discussed the consequences of these findings on user interfaces and proposed a calibration procedure for estimating the useful resolution. Note that our results are specific to the particular mouse we used. Further experiments are required to confirm them with other mice and possibly other devices. For example, results might well be different for touchpads where users can roll and not only move their finger.

ACKNOWLEDGEMENTS

The authors thank H el ene Martin for fruitful discussions related to this work and help with the experiment.

REFERENCES

1. B erard, F., Wang, G., and Cooperstock, J. R. On the limits of the human motor control precision: the search for a device's human resolution. In *Proceedings of INTERACT'11*, Springer-Verlag (2011), 107–122.
2. Casiez, G., and Roussel, N. No more bricolage! Methods and tools to characterize, replicate and compare pointing transfer functions. In *Proceedings of UIST'11*, ACM (2011), 603–614.
3. Chapuis, O., and Dragicevic, P. Effects of motor scale, visual scale, and quantization on small target acquisition difficulty. *ACM ToCHI* 18 (2011), 13:1–13:32.
4. Guiard, Y., Beaudouin-Lafon, M., and Mottet, D. Navigation as multiscale pointing: extending Fitts' model to very high precision tasks. In *Proceedings of CHI'99*, ACM (1999), 450–457.
5. Langolf, G., Chaffin, D., and Foulke, J. An investigation of Fitts' law using a wide range of movement amplitudes. *Journal of Motor Behavior* 8 (1976), 113–128.
6. Meyer, D. E., Smith, J. E. K., Kornblum, S., Abrams, R. A., and Wright, C. E. Speed-accuracy tradeoffs in aimed movements: toward a theory of rapid voluntary action. In *Attention and performance XIII*. Lawrence Erlbaum, 1990, 173–226.
7. Risti c, L. *Sensor technology and devices*. Optoelectronics Library. Artech House, 1994.
8. Roussel, N., Casiez, G., Aceituno, J., and Vogel, D. Giving a hand to the eyes: leveraging input accuracy for subpixel interaction. In *Proceedings of UIST'12* (2012), 351–358.
9. Wobbrock, J. O., Findlater, L., Gergle, D., and Higgins, J. J. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of CHI'11*, ACM (2011), 143–146.