

Toward Goldilocks' Pointing Device: Determining a "just right" Gain Setting for users with Physical Impairments

Heidi Horstmann Koester, Edmund LoPresti
Koester Performance Research
2408 Antietam Dr.
Ann Arbor MI 48105 USA
+1 734 663 4295
hhk@kpronline.com, edlopresti@acm.org

Richard C. Simpson
Human Engineering Research Labs
VA Pittsburgh Healthcare System
Pittsburgh PA 15260 USA
+1 412 383 6593
ris20@pitt.edu

ABSTRACT

We designed and evaluated an agent that recommends a pointing device gain for a given user, with mixed success. 12 participants with physical impairments used the Input Device Agent (IDA), to determine a recommended gain based on their performance over a series of target acquisition trials. IDA recommended a gain other than the Windows default for 9 of 12 subjects. Subsequent performance using the IDA gain showed no meaningful differences as compared to the default setting or users' pre-study settings. Across all gains used by these subjects, however, gain did have a significant effect on throughput, percent of error-free trials, cursor entries, and overshoot. Linear models of gain's effect on performance showed that its effect on throughput is relatively small, with only a 13% difference from highest throughput (at gain = 10) to lowest throughput (at gain = 6). Cursor entries were more strongly affected, showing a steady increase with increasing gain.

Categories and Subject Descriptors

H.5.2 [User Interface]: Subjects – *evaluation, methodology.*

General Terms

Human Factors

Keywords

Computer pointing devices, physical impairment, assistive technology, adaptive computer interfaces

BACKGROUND

Computers have much to offer individuals with disabilities. To fulfill this potential, it is critical that the computer system be closely matched to the user's needs and abilities. The behavior of most computer input devices, such as keyboards and mice, is adjustable. Because each person's disability is unique, tuning these devices to each user's strengths and limitations can be critical for success [1].

Ideally, configuration is performed in consultation with a clinician who has expertise in computer access for people with disabilities. However, a trained clinician may not always be available. Even

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

ASSETS'05, October 9–12, 2005, Baltimore, Maryland, USA.
Copyright 2005 ACM 1-59593-159-7/05/0009...\$5.00.

when a rehabilitation professional is consulted, proper tuning of a device to the needs of a particular user can be a difficult and time-consuming task. The challenge is magnified by the fact that many users' needs and abilities change over time, whether in the short term due to factors such as fatigue or in the long term due to factors such as progression of the individual's underlying impairment. For these reasons, input devices are often not appropriately configured to meet users' needs, with consequent negative effects on user productivity and comfort.

An automated agent on the user's computer could help ensure that input devices are properly configured for the individual, and reconfigured as the user's needs change. Such an agent would need to predict appropriate input device configuration settings based on the user's performance. This requires an understanding of what settings are important, and how input device settings influence performance.

Definition of Gain

One setting which is available for pointing devices is the control-display gain, or sensitivity. Gain determines how far the mouse cursor moves on the screen for a given movement of the pointing device. For the mouse, this is typically measured in pixels (on the screen) per inch (on the physical desktop). The Windows default is 400 pixels per inch. In the Windows control panel, this default setting is assigned a value of "10", and adjustments can be made in the range of 1 to 20.

Effect of Gain on Performance

For mouse users without physical impairments, gain does not appear to have a large effect on pointing performance, at least within a moderate range of the default setting [2]. However, for people who do have physical impairments, from spinal cord injury, neuromuscular disease, cerebral palsy, or other conditions, gain may affect performance, and the "just right" setting may be quite different than the default value. For example, someone with cerebral palsy may have some spastic movements that make it difficult to finely control their pointing device. For this individual, a lower gain might better accommodate their spasticity. Someone with multiple sclerosis, on the other hand, may have good fine motor control but significantly reduced range of motion. For that individual, a higher gain might be a better fit.

In addition to differences between users' physical abilities, there are many other pointing devices (trackballs, joysticks, headpointers, mousekeys, etc.). There are also different movement patterns that can be used with various pointing devices. For example, Zhai [3] has described pointer movements that

follow a constrained path as *steering*, and MacKenzie [4] has referred to *throwing* the cursor across the screen using a single ballistic movement with a trackball. Input methods like MouseKeys and (some) joysticks, on the other hand, allow the user to *drive* the cursor to the target with little or no actual hand movement. Finally, in informal observations we observed several users *crawling* towards the target using numerous small inputs.

The effect of gain on performance might be more important for one or more combinations of device, control method, and user characteristics. The limited literature that addresses this suggests that this might be the case. For example, Radwin and colleagues found that gain did have an effect on movement time for able-bodied individuals using a head pointer [5].

Another study compared subjects' throughput using a head-controlled pointing device for three levels of gain [1]. For subjects without disabilities, there was a moderate but consistent effect of gain. Throughput decreased and overshoot increased as gain increased. However, subjects who had difficulty reaching the targets were better able to do so with increased gain. Among the three subjects with disabilities, the gain resulting in best performance was different for each subject. Bayes' Theorem was used to calculate a decision boundary between those subjects who performed best with the standard interface and those who performed best with the decreased gain. This decision boundary indicated that subjects who had throughput less than 1.7 bits per second or overshoot greater than 10% with the medium gain would experience improved throughput with a decreased gain. On the other hand, subjects who were able to reach less than 90% of the targets would benefit from increased gain.

Adapting Pointing Devices to Users' Needs

If a relationship can be found between an individual's movement patterns and his or her optimal configuration settings, an automated agent could customize the settings in response to the user's needs. McGill developed a force-sensing joystick which adapted to hand tremor using measurements of the user's tracking ability and tremor [6]. A preliminary study with three subjects who had Friedrich's ataxia indicated that the adaptive joystick provided some improvement in performance for tracking tasks. Tracey and Winters developed a system to configure mouse settings in the Windows operating system based on subject performance on computer tracking exercises as well as direct questions [7]. These questions were directed toward an assistive technology clinician who had observed the subject. The system selected appropriate mouse settings which could then be entered manually by the user or a caregiver.

Work has also been done on enhancing pointing performance for people without disabilities. Balakrishnan [2] reviewed a number of techniques for dynamically adapting the pointing device gain or target size based on cursor speed or position; for example, acceleration (higher gain for faster cursor speeds), semantic pointing (lower gain when the cursor is inside targets), and increased target size when the cursor is near targets. While each of these techniques may have some isolated benefits, they may not generalize well to real-world interface use. For movements requiring unusual precision, a cascade neural network with Kalman filtering has been developed which can decrease movement error by an average of 39.5% based on data for microsurgeons [8].

These studies show the potential of pointing device configuration agents for use by people with disabilities. However, they have some limitations:

1. Computer activities often more closely resemble a target acquisition task rather than the tracking tasks used by McGill [6] and Tracey and Winters [7];
2. The system developed by Tracey and Winter [7] assumed the presence of an assistive technology expert; in some cases, users may need or want to configure their systems independently;
3. The system developed by McGill [6] was specialized for a particular group (people with Friedrich's ataxia) and a particular impairment (tremor);
4. The systems described by McGill [6], Balakrishnan [2], and Ang [8] make use of pointing methods that are not available with standard device drivers.

To begin to address these limitations, LoPresti developed and evaluated a system that automatically adjusted the gain for users of head-controlled pointing devices [9]. For 16 subjects with physical disabilities, the system was able to select settings that were appropriate for most subjects and provided a modest but significant improvement in performance ($p < 0.05$). This study provides the foundation for successfully developing a configuration agent for pointing device configuration. However, additional work is needed to apply this work to hand-controlled pointing devices and to provide a greater improvement in performance.

Study Goals

The goal of this study was to determine whether a simple software-based agent, called "IDA", could effectively recommend a specific, fixed gain for users of pointing devices. This agent will adjust gain in a way that is compatible with standard device drivers.

In the process of developing this agent, we wanted to develop a better understanding of how gain affects pointing device performance for users with physical impairments. This is an early step toward developing methods that tailor input devices more closely to users' needs.

METHODS

Subjects

12 individuals whose physical impairments affect their ability to use a mouse took part in this study. All were regular computer users, using their computer at least 10 hours per week, and all had a pointing device that they were able to use, such as a trackball.

Protocol

The protocol began with a recommendation phase, to determine the IDA gain for each subject, followed by an evaluation phase in which each subject's performance was measured under three conditions: (1) "IDA", using the gain recommended by the agent, (2) "Default", using the default gain in the Windows operating system, and (3) "User", using the subject's own gain (which may or may not differ from the default). The order of the three evaluation conditions was counterbalanced evenly across all subjects.

In the recommendation phase, each subject performed 5 runs of a target acquisition task. For the first run, the gain was set to the user's current gain. Based on user performance in each run, the IDA system recommended the gain to use in the next run. At the completion of 5 runs, IDA reviewed the results of all trial sets and made a final gain recommendation.

The target acquisition task used is part of the Compass assessment software package [10]. In each trial, a single target is presented on the screen, and the user moves the mouse cursor inside the target and clicks to select the target. Each test run included 32 trials: 4 targets at each combination of 4 different distances (short, medium, long, and corners) and 2 different sizes (18 and 24 pixels). Short, medium, and long target distances were defined as a portion of display width. Tests were run with a display width of 1024 pixels, yielding small, medium, and long distances of 50, 100, and 512 pixels, respectively. Corner targets were placed in the furthest corner of the screen from the starting cursor position.

Questionnaires

Before testing began, subjects completed a short survey in an interview format. This included demographic information about age, school/work status, education, and type of disability. There were also specific items regarding subjects' use of their computer and pointing device: their experience with it, general pattern of computer use, current gain setting, how that gain was established, and satisfaction with current setup.

Immediately after each evaluation condition, we asked subjects to rate: (1) the gain they just used, on a scale of 1 – 7 (with 1 “way too slow”, 4 “just right”, and 7 “way too fast”) and (2) the task difficulty, also on a 1 – 7 scale (with 1 “extremely easy” and 7 “extremely difficult”). At the completion of all three conditions, we asked subjects to rate the usefulness of a system that helped adjust their gain at times to better match their abilities (with 1 “not at all useful” and 7 “extremely useful”).

Data Analysis

In order to obtain a thorough profile of user pointing performance, 15 different dependent variables were measured. The primary dependent variables of interest in this paper are:

1. Throughput (TP) – throughput is a measure of performance which adjusts target acquisition time according to target distance and width. Specifically, $TP = \log_2(D/W + 1)/t$, where D is the distance from starting position to the target, W is the size of the target, and t is the target acquisition time [1].
2. Percent of error-free trials (EFP) – percentage of the 32 trials that were selected with exactly one click (i.e., with no extraneous off-target clicks);
3. Cursor entries (CE) – the number of times the mouse cursor entered the target during the trial.
4. Overshoot (OS) – Maximum distance traveled beyond the icon as a percentage of the distance to the icon from the starting position.

Performance measures for each trial were averaged across each run of 32 trials with a given gain. The resulting data set had 96 observations: 8 observations for each of 12 subjects.

These data were analyzed as follows to determine whether IDA's recommendations were “effective.” In an absolute sense, EFP in the IDA condition should suggest reasonable competence at the task. There was no fixed a priori criterion established for this, but we felt that an average EFP of 90% would be consistent with competent performance. Additionally, ratings for perceived

cursor speed in the IDA condition should not be significantly different than the “just right” value of 4.

In a comparative sense, the primary question is whether IDA's settings yielded any performance difference relative to the Windows default settings. IDA should certainly not hurt performance, and ideally it would actually help, in a noticeable and clinically significant way. Our rough target for clinical significance was a 20% improvement in throughput. Note that IDA might increase or decrease the gain, relative to the default setting, but the goal in either case is to enhance performance. Differences between IDA and Default conditions were assessed using paired t-tests, with significance defined at $p < 0.05$.

We also conducted follow-up analyses using data from all 8 target acquisition runs to determine the effect of gain on performance. To model the effect of gain while controlling for variation between subjects, we used a linear mixed model with gain as a fixed main effect, and subject as a random effect. A linear mixed model was created for each of the four dependent variables. If a significant main effect of gain was found, pair-wise differences of the marginal means at each level of gain compared to the default gain were estimated, with statistical significance defined as $p < 0.05$. P-values were computed with a Bonferroni adjustment for multiple comparisons.

RESULTS

Participant Characteristics

Participants' age averaged 46 years, and ranged from 27 to 68. Seven were male; 5 were female. Their clinical diagnoses were as follows: 7 with cervical spinal cord injuries; 2 with cerebral palsy; 2 with neuromuscular disease; and 1 with carpal tunnel syndrome. All had graduated from high school, and 9 had at least a bachelors' degree. Nine of 12 had paid employment averaging 36 hours per week, and 4 were in graduate school. Seven subjects used a trackball pointing device; 3 used a typical mouse; 1 used a trackpad, and 1 used MouseKeys to drive the mouse cursor with the cursor keys. All subjects used their own pointing device during the experiment.

Settings in the Evaluation Conditions

As shown in Table 1, IDA changed the gain from the default value for 9 of the 12 subjects, decreasing gain for 3 subjects and increasing it for 6. Four subjects had made their own setting changes prior to this study, and the direction of IDA's changes matched the user's own changes in each case.

Subject	Gain Setting		
	Default	User	IDA
1	10	10	12
2	10	10	10
3	10	6	8
4	10	10	10
5	10	20	20
6	10	10	12
7	10	10	16
8	10	8	6
9	10	10	10
10	10	10	4
11	10	14	12
12	10	10	20

Table 1. Setting values for gain used in the Default, User, and IDA conditions.

Performance in the Evaluation Conditions

Table 2 shows the performance and subjective ratings for each gain condition. In absolute terms, the IDA condition appeared to be readily usable by subjects, with EFP averaging 90.6% and ratings of cursor speed averaging around the neutral value of 4.

However, in a comparative sense, IDA did not enhance pointing device performance. The average performance measures and subjective ratings are nearly identical across conditions, and paired t-tests showed no statistically significant differences. Indeed, looking at each of the subjects individually, only one seemed to benefit noticeably using the IDA setting (Subject 5, with a throughput increase of 14.3%). Conversely, only one subject's performance seemed reduced with the IDA setting (Subject 7, with a throughput decrease of 13.8%).

Measure	Default		User		IDA	
	Mean	95% CI	Mean	95% CI	Mean	95% CI
Thruput (bps)	1.35	(0.93, 1.77)	1.39	(0.94, 1.83)	1.33	(0.91, 1.76)
Error-free (%)	93.2	(83.2, 103.2)	94.0	(89.1, 98.9)	90.6	(81.1, 100.2)
Cursor Entries	1.34	(1.12, 1.57)	1.29	(1.12, 1.46)	1.45	(1.14, 1.76)
Over-shoot (%)	9.0	(5.4, 12.5)	7.3	(3.5, 11.2)	12.5	(5.1, 19.9)
Speed Rating	3.5	(2.6, 4.4)	3.7	(3.3, 4.1)	3.8	(3.0, 4.6)
Difficulty Rating	3.7	(2.6, 4.8)	2.9	(2.1, 3.8)	3.9	(2.9, 4.9)

Table 2. Average performance and subjective ratings across subjects for 3 methods of determining gain. 95% confidence intervals are also shown.

The lack of comparative differences did not change when we looked at the following subsets of subjects: (1) only subjects whose IDA setting was different than the Default (N=9), and (2) only subjects whose IDA setting was at least 4 points different than the Default (N=5).

Main Effect of Gain

One reason why IDA did not yield any significant performance benefits could be that gain was not a strong factor in performance for these subjects. The linear mixed model analyses did show significant main effects of gain for throughput ($p < 0.001$), error-free trials ($p = 0.022$), cursor entries ($p < 0.001$), and overshoot ($p < 0.001$). However, post-hoc analyses of the marginal means suggest that the effect size for throughput is relatively small, although it is statistically significant. (Marginal means are the model estimates for group means at each gain level, after controlling for individual effect of subject.) For throughput, the default gain provides the highest performance, although the gap between the lowest throughput at gain of 6 and the highest throughput at 10 is only 13% (see Figure 1). The only significant difference in throughput, relative to the default value of 10, is for a gain value of 6.

As shown in Figure 1, the cursor entries measure is much more responsive to changes in gain. With the default gain as the reference, cursor entries decrease for all gain values lower than 10, and increase for all gain values higher than 10. There are

statistically significant pair-wise differences at 6, 12, and 16 as compared to 10.

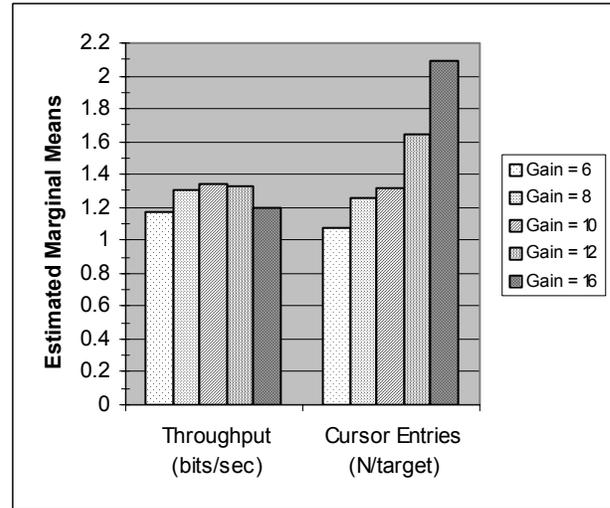


Figure 1. Marginal means at each of 5 gain levels for throughput and cursor entries. Marginal means are model estimates from a linear mixed model with gain as a fixed main effect and individual subject as a random effect.

Test-Retest Reliability

Because all subjects used the IDA setting in the recommendation phase as well as the evaluation phase, we were able to analyze the test-retest variation in pointing performance. The test-retest correlation for throughput with the IDA setting was 0.990 ($p < 0.05$). Somewhat unexpectedly, however, throughput with the IDA setting averaged 6% lower in the evaluation phase as compared to the recommendation phase, a small but statistically significant difference ($p < 0.05$). Conversely, throughput with the Default setting tended to improve slightly in the evaluation phase relative to the recommendation phase, with an average change of 1.9% (not statistically significant). For 4 of the 12 subjects, a notable improvement in Default performance during the evaluation phase, combined with a notable drop in IDA performance, eliminated any advantage of the IDA setting observed in the recommendation phase.

The test-retest “noise” can be estimated by the absolute value of the differences between test and retest throughputs. For the IDA setting, the test-retest noise averaged 7.7%, while for Default, it was 8.4%. Any true effect of gain condition must be detectable above that noise level of approximately plus or minus 10%. This supports our rough guideline for clinical significance as a difference of at least 20% in this domain.

Subjective Results

Regarding subjects' own behavior with respect to adjusting their gain, 9 of 12 said they knew how to change the gain, and 7 of these had actually changed it at one time or another. When asked how the gain on their system had been determined, only one said they had the advice of a clinician with experience in computer access. The rest determined their settings for themselves. Eight of 12 subjects stated that their ability to use their pointing device changed noticeably over time, depending on such factors as fatigue, time of day, and muscle spasticity. Perhaps because of

this fluctuation in ability, subjects had high interest in a system that could adjust gain to meet users' needs, with an average usefulness rating of 5.8 on a scale of 1-7 (95% CI = (4.7, 6.8)).

DISCUSSION

The IDA gain agent made recommendations that seemed reasonable and were readily usable by participants. However, the agent's recommendations didn't provide the performance benefits we were hoping for. This discussion attempts to provide insight into some of the reasons why.

Part of the reason may relate to the specific experimental protocol and participant characteristics. In the recommendation phase, IDA yielded the best performance for each subject, albeit by a small margin in most cases. But this gap closed on the evaluation retest because subjects generally did worse with IDA in the evaluation phase. One reason for this may have been that subjects only experienced the IDA setting for a total of 64 targets, whereas most of them had used the default setting for years. Additionally, only a subset of the participants used an IDA setting that was different from the default setting. This limited our statistical power in examining IDA-Default differences.

Our specific choice of target sizes and distances may have made a difference as well. It is likely that IDA gain changes benefited certain types of targets and harmed others, thus resulting in relatively small net changes in throughput across 32 mixed trials. Preliminary analyses of trial-by-trial data at gain levels of 10 and 12 suggest that this may indeed be the case. Gain of 10 provided better throughput than 12 for all target distances except for corner targets. The difference for corner targets may be due to the fact that, because corner targets are at the edge of the screen, the user is less likely to overshoot; thus the homing requirements are relaxed relative to targets that are away from the display edges. Similarly, gain of 10 yielded better throughput than 12 for small targets (which were about the size of a window close box), but not for toolbar-sized targets. Had we chosen only one type of target size and distance, we might have seen clearer, but perhaps less realistic, results.

Finally, it should be noted that these participants were quite proficient at using their existing pointing devices. They didn't need all that much help, and the kind of help they did need is not the kind that gain adjustments are best able to provide. Behavioral measures such as overshoot and cursor entries didn't reveal notable control problems, which decreases in gain can accommodate, nor did we see problems in reaching all edges of the display, which increases in gain can address [1,9]. While control and range were relatively good for these participants, throughput was less than half that reported for mouse or trackball users without impairments [11]. These subjects mainly needed a boost in throughput, which gain has trouble providing.

This leads to a more fundamental reason why IDA didn't yield notable performance benefits: the effect of gain on throughput is relatively small. Consideration of the equation for throughput helps explain why. An increase in gain reduces the effective target distance, D , by reducing the amount of pointing device motion required to travel a given screen distance. But it also decreases the effective target width, W , by the same principle. The simultaneous changes in effective distance and width tend to cancel each other out. With protocol refinements, we might succeed in demonstrating that a small effect on throughput is

statistically significant, but these results suggest that for many subjects, even those with physical impairments, the effect is unlikely to be of clinical significance. Most of the subjects in this study were people with cervical spinal cord injuries who used a trackball. These data suggest that for these subjects, using gain adjustments as a primary means of accommodating the physical impairment is unlikely to be very productive.

There are other adjustments that might have a bigger effect on pointing performance (e.g., acceleration, and/or adjusting the size of components such as toolbar buttons), and we are beginning to look at those empirically. Additionally, some users may be significantly more efficient using the keyboard, rather than the mouse [12]. For these individuals, time spent tweaking the mouse settings and Windows interface might be better spent in training on how to operate Windows by the keyboard alone. Our goal is to continue to examine ways of configuring the existing interface more appropriately for a user's needs, both initially and over time.

ACKNOWLEDGEMENTS

We thank the participants for their generous contributions of time and effort. This work is funded by a Phase I SBIR grant from the National Institutes of Health, #1 R43 HD045015-01.

REFERENCES

1. LoPresti, E.F., Brienza, D.M., Angelo, J. (2002). Evaluation of Computer Head Control Software to Compensate for Neck Movement Limitations. *Interacting with Computers*. 14(4):359-379.
2. Balakrishnan R. (2004). "Beating" Fitts' law: Virtual enhancements for pointing facilitation. *International Journal of Human-Computer Studies*, 61(6):857-874.
3. Zhai, S., Woltjer, R., (2003). Human Movement Performance in Relation to Path Constraint - The Law of Steering in Locomotion, in *Proceedings of IEEE Virtual Reality 2003*, pp 149-156.
4. MacKenzie, I.S. (1992). Fitts' Law as a Research and Design Tool in Human-Computer Interaction. *Human-Computer Interaction*. 7:91-139.
5. Radwin RG, Vanderheiden GC, Lin M-L. (1990). A method of evaluating head-controlled computer input devices using Fitts' law. *Human Factors*, 32(4):423-438.
6. McGill, R. (1990). A Blackboard Knowledge-Based Approach Towards Implementing an Adaptive Force Joystick Computer Input Device for Persons with Tremor Disability. *Proceedings of the RESNA 1990 Annual Conference*. p. 431-432.
7. Tracey, M., Winters, J. (1999). Neuro-Fuzzy Advisor for Mouse Setting in Microsoft Windows. *Proceedings of the BMES-EMBS 1st Joint Conference*. p. 664.
8. Ang, W.T., and Riviere, C.N. (2001). Neural Network Methods for Error Canceling in Human-Machine Manipulation. *Proceedings of the IEEE Engineering in Medicine and Biology Society Conference*.
9. LoPresti, E.F., Brienza, D.M. (2004). Adaptive Software for Head-Operated Computer Controls. *IEEE Transactions on Rehabilitation Engineering*. 12(1):102-111.

10. Koester, H.H., LoPresti E.F. (2003). Compass: Software for Computer Skills Assessment. *CSUN's 18th Annual Conference "Technology and Persons with Disabilities"*. March 2003.
11. MacKenzie, I.S., Kauppinen, T., Silfverberg, M. (2001). Accuracy Measures for Evaluating Computer Pointing Devices. *Proceedings of the ACM Conference on Human Factors in Computing Systems – CHI 2001*, pp. 9-16. New York: ACM.12.
12. Cantor, A. (1997). An Evaluation of Keyboard-only Access to Windows for "Single-Digit" Typists. *Proceedings of the RESNA 1997 Annual Conference*, pp. 357-359.