

Automatic Assessment and Adaptation to Real World Pointing Performance

Amy Hurst

Human Computer Interaction Institute, Carnegie Mellon

akhurst@cs.cmu.edu

Abstract

There is a growing population of individuals who are motivated to use a computer but find it physically difficult to do so, particularly when using a pointing device such as a mouse. Common pointing problems include inability to select small targets, difficulty moving a pointing device in a straight line, or difficulty controlling the pointer's buttons. There are many software adaptations that can improve pointing performance, however the key to these solutions is identifying when to deploy them. In my thesis, I am working to build software that can automatically assess pointing problems and their severity during real world computer use, and deploy appropriate assistive adaptations. In order to evaluate real world interaction with GUIs, I am developing both data capture and analysis tools that are application independent. This software-based approach to improving computer access is powerful because it will be able to adapt to an individual's changing needs and work across applications.

Introduction

One of the main reasons computers are inaccessible is that they treat all users the same, and usually do not accommodate a user's changing needs. Currently the best solution to match ability to technology is expensive human-in-the-loop technology assessment. Unfortunately, the high cost of these assessments prevents them from being performed frequently enough to capture an individual's changing ability. As a result, it is not uncommon to find individuals with pointing problems using either no computer access technology or technology that is not configured for their needs.

I am working to build software that can automatically assess a user's pointing performance during their computer use and adapt to their current needs. By assessing performance during real world use, assessments can be made more frequently and can better model a user's changing needs.

My thesis works towards this goal by focusing on three aspects of this problem: real world data collection, assessment and adaptation. In order to maximize the potential impact, I am working to make all pieces application independent and, when possible, operating system independent. I am collaborating with individuals with a wide range of motor and cognitive abilities at United Cerebral Palsy Pittsburgh to observe real world pointing problems and collect both laboratory and real world pointing data.

In this article I will briefly review related work on pointing performance and adaptations designed to accommodate pointing problems. Then I will discuss my progress collecting a dataset of real world pointing use from individuals with and without motor impairments, building learned statistical models to assess pointing performance and deploying pointing adaptations to support user needs.

Related Work

Pointing Performance

Many researchers have studied pointing performance trends of individuals with a wide range of abilities in laboratory settings including [1,3,9,10,11]. These studies have characterized pointing performance in terms of movement trajectories, accuracy, clicking behaviors and speed. I draw from these studies to better understand pointing performance, common pointing problems, and how to assess them.

While these studies have answered many questions about input devices, interface designs and human ability, they have answered many fewer questions about what happens outside the lab, and if these findings transfer to real world use. My work strives to learn more about the characteristics of real world pointing use, compare it to laboratory data by collecting real world pointing data, and analyze it with the same techniques used in laboratory studies.

Interface Adaptations

There is a large body of work exploring how to optimize interfaces for performance and make pointing tasks easier. Examples include changing the way the cursor interacts with targets to make them easier to access [12], increasing or decreasing pointing gain [10], and pausing the pointing device during a click to prevent it from slipping off a target [11]. These adaptations are all designed to accommodate specific pointing problems, and are able to improve performance for individuals who have those problems. However, the key to a large-scale successful deployment of these adaptations is in knowing when to deploy these adaptations and the optimal settings for the current user.

An alternative to changing how the pointing device interacts is to change the onscreen interactors to meet the needs of the user or the input device. One solution to doing this includes collecting both preference information and baseline performance data from a user and automatically designing an optimized GUI for their needs [4]. Another solution is to automatically tailor the interactors in an application to the input devices available to the user. Configurations include redesigning an interface for interaction with a keyboard and mouse, only a keyboard, only a mouse, or even a single switch button [2]. While these approaches represent a valuable way to improve accessibility, these techniques are limited in the applications and toolkits that they can support.

Assessment and Adaptation

Data Collection and Analysis

The best way to build models that will generalize is to use large datasets of pointing actions that represent a wide range of problems. Ideally these examples are representative of problems that happen infrequently. One relatively underexplored area of pointing performance is the study of real world pointing performance. I am exploring this in a study to collect real world pointing performance from individuals without pointing problems. This dataset will be used as a baseline of real world performance data to compare to data from individuals with pointing problems. I am currently conducting a long-term study collecting data from individuals with cognitive and motor impairments at United Cerebral Palsy Pittsburgh.

I have conducted an initial analysis comparing laboratory trials to real world pointing data that found high variance within individuals over time [7]. Figure 1 illustrates the high variance in distance slipped while clicking a target across different login sessions of real world use (sessions 1 through 9) and compared to a baseline laboratory pointing task (session 0). This finding helps motivate the goals of this thesis because one of the best ways to support high variance is through frequent assessment.




Figure 1: Real world pointing data from 3 users, illustrating their variance in pixel distance slipped during a laboratory evaluation (session 0) and across real world use (sessions 1 and above).

One of the largest challenges in analyzing real world pointing data is automatically determining the sizes of the targets interacted with. While accessibility APIs such as the Microsoft Active Accessibility API (MSAA) provide information about some of the targets, they leave many unsupported. I have developed a technique to augment information from Accessibility APIs with computer vision to improve automatic target identification in GUIs. In a dataset of 438 images from 6 applications, my hybrid target identification technique improved overall recognition accuracy from 18.5% (with the MSAA alone) to 73.5% (with MSAA and computer vision). This technique finds targets by analyzing the visual difference before and after a button press, or by using template matching of common target edges to identify a target. Figure 2 illustrates how the MSAA API was unable to identify the target, but my technique found the target successfully. With this dataset I will also investigate the distribution of pointing types and characteristics of frequently used

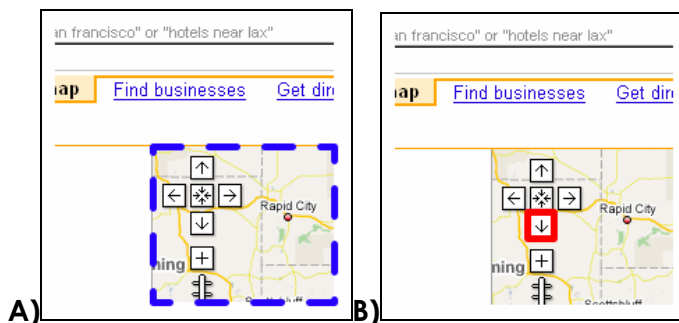


Figure 2: Target selection technique identifying targets in maps.google.com A) Dashed rectangle indicates smallest target MSAA API can detect. B) Bold square indicates target user pressed, which was found with the template matching part of my technique.

targets. Specifically, I am interested in how well real world movements actually fit pre-existing performance models such as Fitts' Law [3] and the Steering Law [1]. I will also explore what categories of motion our users perform which cannot be described by these models. This dataset will enable me to investigate the distribution of target sizes and the distances between them, and see how representative they are of the targets and distances used in laboratory studies. Most importantly, I will use this dataset as the training set to build learned statistical models that are able to assess real world pointing.

Assessing Pointing Performance

I have built learned statistical models of pointing performance to automatically detect characteristic movement from individuals with a wide range of abilities. These models distinguish between different groups of use with performance metrics, or features, that describe a pointing interaction. To date, I have built models assessing pointing performance from laboratory studies evaluating menu use and 2D target acquisition.

To better assess performance while using menus, I collected a dataset of novice and skilled menu selections made while using an image editing program. The interaction between the pointing device and menus was analyzed and scored to develop a set of performance metrics, or features, for each menu selection which were used to build a learned statistical model. This model was able to distinguish between novice and skilled menu use with 91% accuracy, without using a task model of the user's actions [5].

To better assess performance of 2D target acquisition, I built learned statistical models with 3 datasets from prior work in this area to distinguish between performance differences due to age and motor impairment [6]. Specifically, data from the Koester dataset [10] was used to distinguish between the pointing performance of individuals with and without known pointing problems. A dataset from Keates [9] was used to distinguish between pointing performance based on age groups or the diagnosis of Parkinson's Disease. Finally, a dataset from Trewin [11] was used to predict how well a learned statistical model could predict if the Steady Clicks adaptation would improve performance of individuals who slipped off a target when they tried to click. All 3 models performed with high accuracy results (above 90%), and are illustrated in Figure 3.

Adapt to Pointing Performance

The learned statistical models described in the previous system can be incorporated into applications so they can continuously assess a user's performance and act accordingly. To adapt to pointing performance, an application should deploy an adaptation designed for a specific performance change once it detects that such a change has occurred. Depending on the specific pointing problem, this adaptation might change the way the pointing device interacts with onscreen elements, or the way it responds to user movement.

I am currently working to build a suite of tools to log user actions and plug-and-play pre-existing adaptations. I have developed one solution to help users acquire frequently used targets by observing a user's click and drag history [8]. With this history, my software applies a simulated gravity field around the collection of previous interactions to help the user acquire targets. This technique is application independent and operating system independent. While this technique provides information about regions that contain targets, it is limited to adaptations that are agnostic of target size because it does not

know their precise size or exact location. I have developed the hybrid computer vision and MSAA API technique to more accurately identify target size and location so more adaptations can be supported.

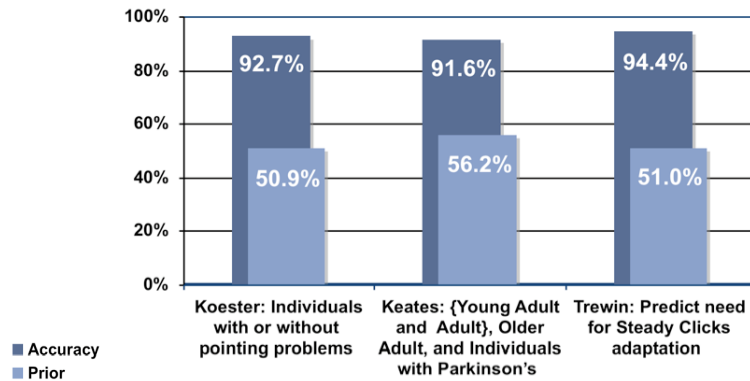


Figure 3: Performance of three learned statistical models built from three datasets from other researchers. These models distinguish pointing performance to predict presence of pointing problem, age differences, and if an individual's performance would improve with Steady Clicks adaptation.

Status of Research

Collect and Analyze Real World Data

I am working to collect a large dataset of pointing performance from individuals with a wide range of physical abilities performing real world tasks. This real world dataset will be unique and will be analyzed to learn more characteristics of real world pointing including how technique, accuracy, and pointing problems differ over time and across users. In order to perform this analysis, I will employ my techniques to identify onscreen targets using the MSAA API and computer vision. Upon completion of my thesis, I plan on making this database available to other researchers interested in studying real world GUI interactions.

Assess

I plan to use this real world dataset to build learned statistical models to identify computer users who would benefit from accessibility tools and make adaptation recommendations. Based on the performance of these models, I will be able to design an appropriate system to translate the output of a classifier into specific adaptations and changes that need to be made.

Adapt

Once I have built models based on real world pointing behavior, I will evaluate the performance of these models and automatically deploy adaptations to the interface that are tailored to the user's current ability. Participants will be interviewed about their perceived performance when the adaptations were deployed and how the adaptations affected their pointing performance. These subjective results will be compared to their performance with and without the deployed pointing adaptations.

Conclusion

I am working to make computers more accessible by automatically assessing pointing performance and adapting interfaces to address pointing problems. I am collecting a database of real world pointing behavior, and have had success analyzing real world pointing data, building learned statistical models to identify pointing problems in a laboratory study, and building adaptive systems. With this research, I hope to make computers more accessible for individuals with pointing problems.

Acknowledgements

This work was supported in part by the National Science Foundation under grants EEC-0540865, IIS-0713509, IIS-0325351, Amy Hurst's NSF Graduate Student Research Fellowship; IBM Research; and the Pennsylvania Infrastructure Technology Alliance.

References

1. Accot, J., & Zhai, S. (1997). Beyond Fitts' law: models for trajectory-based HCI tasks. *Proceedings of the SIGCHI Conf. on Human factors in computing systems*, 295-302.
2. Carter, S., Hurst, A., Mankoff, J., & Li, J. (2006). Dynamically adapting GUIs to diverse input devices. *Proceedings of the 8th international ACM SIGACCESS Conf. on Computers and accessibility*, 63-70.
3. Fitts, P. (1954). The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology*, 47, 381-391.
4. Gajos, K., Wobbrock, J., & Weld, D. (2008). Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces. *Proceedings of the SIGCHI Conf. on Human factors in computing systems*, 1257-1266.
5. Hurst, A., Hudson, S., & Mankoff, J. (2007). Dynamic detection of novice vs. skilled use without a task model. *Proceedings of the SIGCHI Conf. on Human factors in computing systems*, 271-280.
6. Hurst, A., Hudson, S., Mankoff, J., & Trewin, S. (2008). Automatically detecting pointing performance. *Proceedings of the international Conf. on Intelligent user interfaces*, 11-19.
7. Hurst, A., Mankoff, J., & Hudson, S. (2008). Understanding pointing problems in real world computing environments. *Proceedings of the ACM SIGACCESS Conf. on Computers and Accessibility*, 43-50.
8. Hurst, A., Mankoff, J., Dey, A., & Hudson, S. (2007). Dirty desktops: using a patina of magnetic mouse dust to make common interactor targets easier to select. *Proceedings of the ACM Symposium on User Interface Software and Technology*, 183-186.
9. Keates, S., & Trewin, S. (2005). Effect of age and Parkinson's disease on cursor positioning using a mouse. *Proceedings of the ACM SIGACCESS Conf. on Computers and Accessibility*, 68-75.
10. Koester, H., LoPresti, E., & Simpson, R. (2005). Toward Goldilocks' pointing device: determining a "just right" gain setting for users with physical impairments. *Proceedings of the ACM SIGACCESS Conf. on Computers and Accessibility*, 84-89.
11. Trewin, S., Keates, S., & Moffatt, K. (2006). Developing steady clicks: a method of cursor assistance for people with motor impairments. *Proceedings of the ACM SIGACCESS Conf. on Computers and Accessibility*, 26-33.
12. Worden, A., Walker, N., Bharat, K., & Hudson, S. (1997). Making computers easier for older adults to use: area cursors and sticky icons. *Proceedings of the SIGCHI Conf. on Human factors in computing systems*, 266-271.

About the author:



Amy Hurst is a PhD student at the Human Computer Interaction Institute at Carnegie Mellon in Pittsburgh, Pennsylvania where she is advised by Scott E. Hudson and Jennifer Mankoff. Her research interests include Assistive Technology and Context Aware Computing. She graduated from Georgia Institute of Technology in 2003 with a BS in Computer Science and a Certificate in Film Studies. She is the recipient of a National Science Foundation Graduate Research Fellowship (2005) and an Anita Borg Scholar (2008).