In D. Reitter & F. E. Ritter (Eds.), Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016). University Park, PA: Penn State.

Using Naturalistic Typing to Update Architecture Typing Constants

Marc T. Burns (mtb137@psu.edu)
Frank E. Ritter (frank.ritter@psu.edu)
Xiaolong "Luke" Zhang (lzhang@ist.psu.edu)
College of Information Sciences and Technology, Penn State University
University Park, PA 16802 USA

Abstract

Despite the near-ubiquity of graphic user interfaces for navigating the digital and virtual space, relatively little is known about their naturalistic usage. We start to address two questions. First, how do people use computers outside of laboratory studies? This includes what we can we determine about user behavior by analyzing detailed user logs. Second, can we update the constants in user and cognitive models for predicting typing time based on naturalistic behavior? We thus recorded naturalistic logs of 45 users over 219 sessions providing 1,865 hours of behavior (average session=8.18 hours). The analyses of keystroke times are sensitive to the definition of typing (e.g., how close keys are to be counted as continuous typing), and comparisons will need to provide clear definitions or tradeoff curves. Using this data, we updated the typing and homing constants of the Keystroke-Level Model (Card, Moran, & Newell, 1983), a theory of interface behavior used to provide constants for many cognitive architectures. The results suggest that people are typing faster than previously believed based on the 1983 KLM predictions; homing (moving one's hand from mouse to keyboard and keyboard to mouse) occur frequently, and now appear to be different events and thus require separate constants.

Keywords: GOMS, KLM, Typing Rates

Introduction

How do people use keyboards when they are not working in the lab? To investigate this, we performed a naturalistic study of how users interact with their computers in their natural settings. Because users are spending increasing time interacting with their computers earlier in life, the need to understand how users interact is increasingly important.

This study differs from previous research by analyzing the results from users' natural computer usage. By studying users in a variety of settings where they use their own computers, in participants' offices and homes, we are able to log users' actual behaviors as they occur. This behavior differs from previous laboratory studies due to the naturalistic approach to this study.

Studies focusing on how users perform tasks in labs risk diminished authenticity of the data collected because users may be working on unfamiliar computers in an unfamiliar context (the laboratory setting) using an unfamiliar keyboard to do a new, prescribed task. For example, few if any studies in our experience ask participants to set the keyboard settings to their preference. This study addresses these issues through a naturalistic study of human-computer interaction (HCI).

Users' interaction with their computer interfaces is significant for several reasons. First, one can develop an understanding of users' inefficiencies and capabilities. Second, targeting users' fundamental inefficiencies and building from their capabilities can provide developers insight into designing easier and more efficient interfaces. It may be necessary to study users' natural interaction tendencies to predict accurate interface times.

Lastly, a remaining challenge facing the study of users, and more broadly HCI, is that of accurately mapping what users actually do, when they do it, and why they do it. Lab studies like Card et al. (1983) can help us gain accurate understanding of user interaction behaviors, but this method assigns tasks to users, rather than lets users organically select their own tasks. While their experimental design was highly effective in quantifying specific tasks, it is probably inaccurate to assume that modern computer users follow the same single tasked-ness (e.g., Salvucci & Taatgen, 2011; Spink, Ozmutlu, & Ozmutlu, 2002). Additionally, experiment-driven research cannot account for patterns of user task selection, task switching behavior, and users' breaks in natural computer use. For example, few experiments account for users watching videos during the completion of their task. Thus, to gain an accurate representation of how users act in their day-to-day usage, naturalistic research is required. Furthermore, little research has been completed since Card, et al. (1983) to update the constants in the KLM; these may have changed with increased usage and with a broader population of users.

Previous Studies and the KLM Model

We note here several studies that make suggestions about how we should proceed with this work.

Early Studies and the KLM

Kinkead (1975) investigated the differences in keyboard layouts, and found that there is a nominal difference across layouts, suggesting that the speed of typing is based more on how practiced a keyboardist is, rather than an "optimal" keyboard layout. Kinkead concluded that a vast majority, 95%, of keystrokes occur in 2/3 s, or 667 ms, which is much slower than the times defined later in Card et al. (1980). However, similar to Card el al., these users performed prescribed tasks.

In 1980, Card et al. published a paper outlining the Keystroke-Level Model (KLM), a simple tool to aid in the designing of interactive computer systems. The KLM used the time required to perform the sub-steps that make up a task on a computer to predict the time it takes an expert to perform said task (Card et al., 1980). The original purpose of the KLM was a rough calculation for system design. However, in creating the KLM, Card et al. extensively studied the sub-tasks, including physical-motor operators that make up the task, as well as the mental priming required to perform the task.

Keystroking is the time required to strike one key. The KLM treats all keystrokes the same, regardless of the key. These times range from 0.4 s per character for the best typist to 1.20 s per character for the worst typist on an unfamiliar keyboard.

Homing time is defined as the time for a user to move their hand between devices. This motion is from the keyboard to the mouse, or vice-versa. Card et al. reported homing time as 0.4 s during a text-editing task.

Surprisingly, despite dramatic changes in the early adoption, frequency of use and format of computers, we find few recent studies in this area. Additionally, even at the time of Card et al., the rate-limiting step of computer users in text-editing tasks was not the interface of their computer, but rather the information-processing capacity of the user (Card & Moran, 1986).

The first challenge is to expand the applicability of the research by Card et al. (1980). To simplify the process of creating models that accurately describe computer user behavior, Card et al. solely studied the task of text editing under several strict restrictions:

- Users could only be expert computer users
- The task must be routine
- Performance must be without errors.

This contextual rigidity allowed Card et al. to develop the Keystroke-Level Model and GOMS (Goals, Operators, Methods, and Selection rules), two widely used human-information processor models. However, these models do not apply to all users, tasks, and metrics. Both the Keystroke-Level Model and GOMS focus on time required by experts to complete a task as the unit of measurement. These constants are used in ACT-R and EPIC, and sometimes with Soar models to check timing predictions. These models are unable to address other fundamental metrics such as the quality of work. They also did not study users doing tasks that users complete outside of lab experiments.

Thus, this paper focuses on the physical-motor tasks of typing and homing, and the time required to perform these tasks. Despite the KLM model allowing 1.35 s mental preparation at the beginning and ending of tasks, as well as system response time, which is dependent upon the system being used, our study examines what constitutes continuous typing, the time to perform a keystroke within continuous typing, what comprises of a homing, and the time to complete a home.

Card and Moran (1986)

Three years after the publication of *The Psychology of Human-Computer Interaction*, Card and Moran (1986) revisited some of the assumptions and building blocks upon which their book was written. They outlined four interfaces for modeling the interaction between humans and computers: physical, cognitive, conceptual, and task. These four interfaces constitute the foundation of the literature that this paper updates.

Later Updates

Since the studies of Card, Moran, and Newell, many researchers have moved to answer remaining questions, on how users interact with computers. Unfortunately, studies to answer these questions typically and nearly exclusively occur in carefully controlled, task-oriented contexts (MacKenzie, Sellen, & Buxton, 1991; Whisenand & Emurian, 1999).

To that end, we study participants in a naturalistic environment to obtain data more representative of actual behavior and to build a comprehensive representation of how GUIs are used in daily computer interactions.

Summary

Thus, the constants in the KLM theory have been used for a while and could be updated based on users' increased experience with interfaces. The data used in the KLM have been obtained from experimental studies, as opposed to taken from users performing their own tasks on their own devices in a naturalistic setting. To improve our understanding, we will record users' mouse movements and keystrokes while they perform their normal daily tasks on their own computers.

Method

Data was collected by recording the users' keystrokes and mouse movements while they performed their typical tasks at home and at work. We used an anonymizing keystroke logger for privacy.

With this data, we start to determine what constitutes typing, and the naturalistic typing speed; a rich area that we are just beginning to explore. We also provide an update to the typing constants from Card et al. (1983).

The input logger records mouse clicks, mouse movements, and keystrokes across all tasks. Furthermore, such information is useful for inferring patterns of natural behavior, such as how long participants use a computer at a session, total daily usage, number of keystrokes, typing speed, etc.

Participants encounter the logger only upon the start and conclusion of a session, without artificial tasks, distractors, or external observers or apparatus.

Participants

Of the 45 unpaid participants 18 were male, 12 female, and 15 declined to report. In addition, 3 participants were left-handed, 17 were right-handed, and 25 declined to report handedness; however, each participant used their preferred hand for using the mouse. Participants were members of the Penn State community, with a majority (20) being undergraduate students in the College of IST, 16 employees including office staff and professors, and 9 graduate students. We assigned subjects IDs between 1 and 100.

Apparatus

Participants used their own hardware: 44 used Windows on PCs, while one used OS-X on Apple hardware.

We recorded the users' inputs with Recording User Input ver. 2.03 (RUI), a keystroke and mouse logging tool (Kukreja, Stevenson, & Ritter, 2006; Morgan, Cheng, Pike, & Ritter, 2013) that runs on Windows and Macintosh computers. This modified version of RUI anonymized most keystrokes. Individual keystrokes were replaced with the asterisk character; combination key keystrokes (chords), including shift, alt, and control preserved the special key being pressed with an asterisk to signify a keystroke.

Design and Procedure

Participants were provided a thumb drive with RUI or had RUI installed on their computer by an experimenter. RUI ran in the background while participants used their computers. Sessions typically lasted a workday (mean of 8.18 hours), ending when the participant terminated RUI and the thumb drive was returned to the experimenter. Participants typically recorded four sessions.

Keystrokes included as continuous typing in later analyses met the criteria in Table 1, while user actions included as homing in further analysis met the criteria defined in Table 2.

Table 1: Criteria for including keystrokes in continuous typing calculations.

- Contiguous (no intervening mouse events)
- Not the first of its action, e.g. must follow a keystroke
- Not an alt or control key chord
- Occur within 2 s of the previous keystroke

Table 2: Criteria for Homing.

- Mouse movement following a keystroke; or keystroke followed by a mouse movement
- An additional mouse movement occurs following the initial mouse movement of 10 or more pixels; or an additional keystroke must follow the initial keystroke
- Occurred in 2 s or less from previous action

Results and Discussions

Overall, our 45 participants logged 1,816 hours of interaction including over 1.5 million keystrokes. From the ~60.3 million records (over 3 GB), we computed derived measures that we present in two sections: (a) the naturalistic typing behaviors, (b) homing actions.

Analysis of Keystroke Data

What is continuous typing? The typing rate can be calculated as the total number of keys divided by time for a fixed amount of text. Analysis of the raw data showed an average typing speed over the whole sessions (including all pauses) of 13.7 characters per minute or 2.25 words per minute (wpm) assuming 82% of characters typed are word characters with a word of 5 characters (MacKenzie, 2002; Salthouse, 1984).

Because our data is naturalistic, it is not possible to know when the intention to type starts. So we needed to determine a cutoff for when typing begins and ends. In our analysis, we define typing as two or more contiguous keystrokes non-interrupted by mouse movements or button clicks. We computed and plotted the mean time between keystrokes. Figure 1 shows how the typing rate varies as the time threshold allowed between keystrokes varies from 1 ms to 30,000 ms across all users. At 0 ms, one cannot compute a typing rate, and 30 s represents a relatively long time between keystrokes.

In Figure 1, at around 10 s, the mean keystroke time starts to flatten. Ten seconds is still probably too large to consider being a dwell time between keystrokes in continuous typing. Card, Moran, and Newell (1983) in the KLM and Kieras in GOMS (Kieras, 1988) note 1.20 s per keystroke for an inexperienced typist on unfamiliar devices. We thus round up and use 2 s as a threshold in the remainder of this paper, although Figure 1 shows that the computation of keystroke time is sensitive to this threshold.

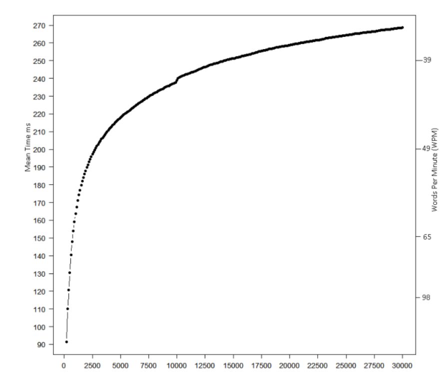

Figure 1: Mean time for keystroke (ms) vs. cutoff time (ms).

Figure 2a shows that the distribution of the times between keystrokes (with < 2 s separation) has a long tail. Figure 2b shows that individual distributions also have long tails. This effect is likely caused by distractions, such as leaving the keyboard, other task actions performed on other devices such as phones, or other events occurring while the user was typing or entering short strings.

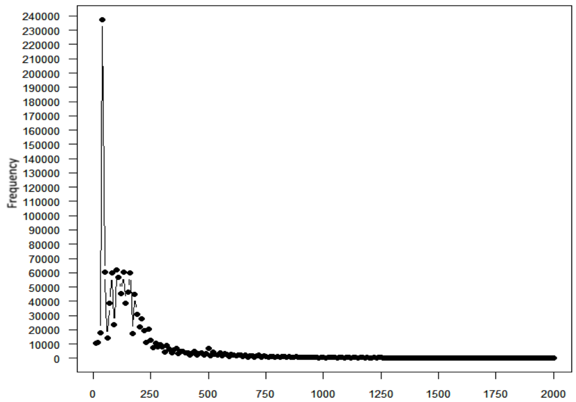

Figure 2a: Keystroke time distribution 2 s threshold in 2 ms bins across all users.

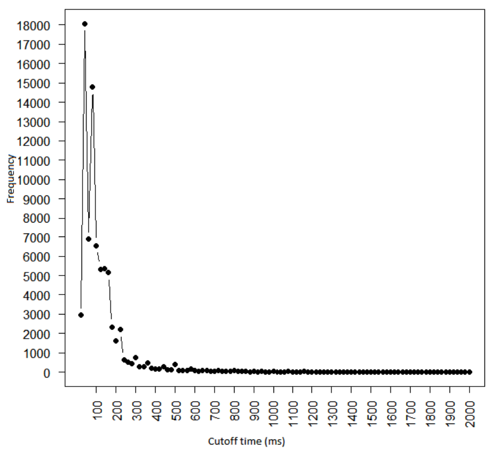

Figure 2b: Subject 6's keystroke distribution with 2 s cutoff in 20 ms bins.

Using a 2 s threshold (arbitrarily chosen based on examining Figures 1 and 2), we observed that the overall mean keystroke time is around 195 ms, which falls between average (200 ms) and good (120 ms) typists. Some latency in keystroke times may be attributed to interruptions or other externalities not controlled in this environment.

Keystrokes

Figure 3 shows the distribution of the 1.512 million keystrokes by users meeting the criteria in Table 1 (95.6% meet the criteria). It shows that users type different amounts. Table 3 shows the relationship of dwell time between keystrokes and the words per minute. The mean keystroke time was 195.5 ms, equal to 296.7 characters per min. and 48.3 wpm. The median keystroke time was 191.2 ms. Figure 4 shows the distribution of keystroke times per user, with a minimum average time of 76.5 ms (128.5 wpm) and a maximum of 291.7 ms (33.7 wpm). For comparison, the base keystroke time of the KLM is 200 ms (49.2 wpm) for secretaries (professional) and 280 ms (35.1 wpm) for non-secretary (non-professional) typists (Card et al., 1983). Our findings suggest that users type faster than their cohorts in 1983.

Table 3: Relationship of keystroke times to words per minute (wpm)

Keystroke Time (ms)	Keys / Minute	Word Chars / Minute	Words / Minute
50	1200	984	197
75	800	656	131
100	600	492	98
150	400	328	65
200	300	246	49
250	240	197	39
300	200	164	33
350	171	141	28

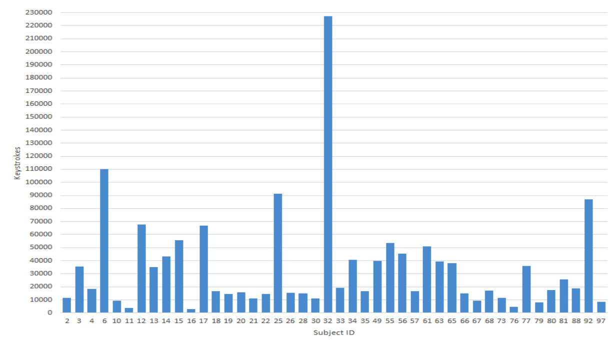

Figure 3: Distribution of keystrokes by user using 2-s cutoff. X-axis labels indicate participant IDs.

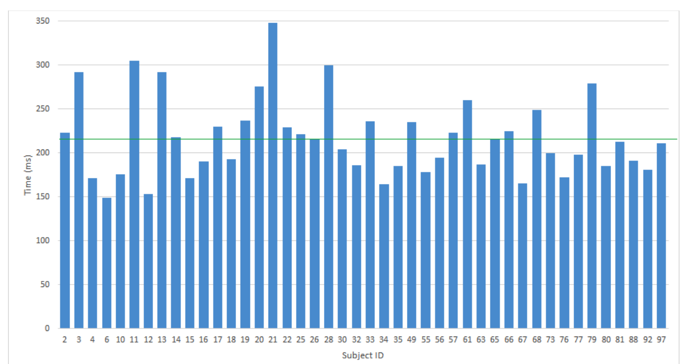

Figure 4: Distribution of mean keystroke times by user.

Homing

Users homed (moved their hand) from keyboard to mouse 51,989 times, and from mouse to keyboard 45,302 times (total: 97,291 times). Table 2 shows the criteria constituting homing. Combined, this rate is 53.6/hour or 0.89/minute. This action occurs once for every 15.4 keystrokes (using all keystrokes).

Figure 5a and 5b shows the distribution of homing times between the keyboard to mouse and vice-versa. The time to home took less time in either direction than the 0.40 s predicted by the KLM (Card et al., 1983). Homing had a mean value of 0.16 s for keyboard-to-mouse and 0.32 s for mouse-to-keyboard. The combined mean is 0.230 s. The difference in the times between directions implies that the user needs additional time to return hands to the keyboard to continue typing; while when starting to mouse, the user may often be still typing while the hand is moving towards the mouse. These results suggest that the direction of homing (mouse to keyboard vs keyboard to mouse) are different actions, do not occur equally often, and could be separated into two distinct constants to give more accurate predictions.

We also observed a large number of occurrences—8051—of 0 ms homings (combined). These times may indicate co-joint typing and mousing or typing, or accidental trackpad movements.

Figure 5b and to a lesser extent Figure 5a, show a small peak in the distribution around 500-700 ms. The second peaks of the distribution were not expected; their existence might be due to two different user behaviors. We speculate that the 0-100 ms response time represents homing when one hand is on the keyboard and the other is simultaneously on the mouse, allowing near instantaneous homing, while the other response time represents a short delay between keystroke and homing.

Therefore, we see that movements between keyboard and mouse occur about once for every 15.4 keystrokes, and that this time is likely now done in two strategies: (a) while typing one moves the mouse with the other hand, and (b) stop typing to use the mouse with the last hand to type. The former is somewhat faster. We also saw that time to keyboard and time to mouse might be different actions. It might be useful to separate these two actions into separate actions to improve predictions where this is desirable.

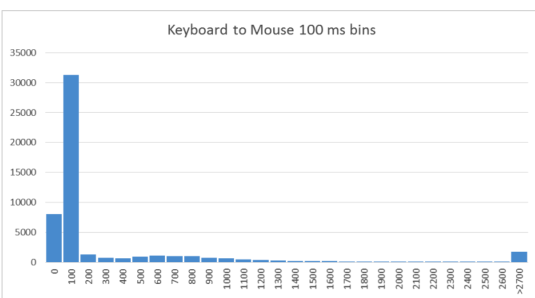

Figure 5a: Distribution of keyboard to mouse times; mean of the whole distrbution is 0.16 s.

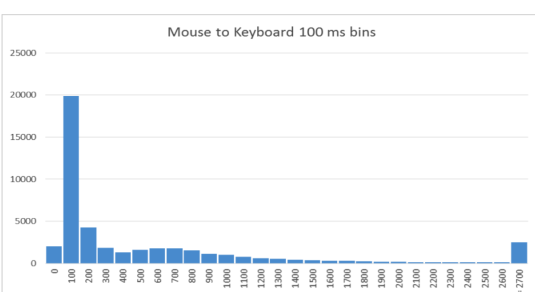

Figure 5b: Distribution of mouse to keyboard times; mean of the whole distrbution is 0.32 s.

We also saw that from the relative frequency of the two types of homing behaviors, that users tend to end a short set of unit tasks on the keyboard (there are more mouse to keyboard actions) than end a session with using the mouse (there are fewer keyboard to mouse actions).

There can be some ambiguity to what constitues homing, as the user may have their hand on the mouse causing a stream of mouse and keyboard data that still adheres to the criteria we proposed. Knowledge of the keys typed as well as the preferred mousing hand of the user could lessen this ambiguity. Furthermore, if the user is using a laptop with built in trackpad, the user could be mousing with their thumb and the remaining fingers not leaving the keys.

Conclusions

This study can only be a start to exploring naturalistic behavior with computers—there is a much greater range of environments and types of users than we could cover here. However, we were able to provide a general summary of naturalistic user behavior and to update several user constants used by many cognitive architectures. Naturalistic user studies of computer interaction can provide a new domain for large data analyses.

We found the definitions of these measures more important than we initially thought. For example, the definition of typing or homing depends on more than just dwell time. There were many additional hidden aspects. For instance, how do we detect when a user was away from the keyboard? When did the user switch from the keyboard to the mouse? Was the user moving the mouse and typing at the same time or was there desk instability that lead to the mouse appearing to have moved whilst the user was typing?

Updating the KLM Constants

We saw in this naturalistic study that the typing rate is sensitive to typing threshold. We can use this approach in later analyses. This result also reminds us that typing speed is not a fixed rate but a distribution, and that the tails of this distribution might be important for some analyses.

The amount of typing over a workday is not very large, (average of about 7,700 keystrokes). Subsequently, the typing rate over a workday is not fast (about 2 wpm). While the amount of typing varied by user by over a factor of 20, the typing rate varied by only a factor of 2.

These results suggest that users are typing faster than Card et al.'s (1980) KLM expert times, and that current naturalistic typing is not much slower than previous timed typing.

We suggest that architectures that use these constants (e.g., ACT-R, EPIC) (a) update their typing rate, (b) allow for additional concurrent motor commands, (c) split homing actions into two distinct actions, and (d) use these shorter times for homing tasks.

Limitations

While we include a range of faculty, staff, and students, our participants were all from the same university environment.

A wider selection of users is likely to provide a wider range of behavior. Users from different work domains, a greater number of users using computers more at home or in other environments, and older and younger users might lead to a different constants and different results. It may be useful to rerun this study with those more varied types of users.

While we analyzed a large amount of data, observing additional users could provide more support for these conclusions.

We were limited by the anonymization of the data: the keystroke logger anonymized many of the details of the data making it difficult to determine what, task the users were performing. While anonymization was necessary for the protection of participants' privacy, with non-character key data such as return, arrow keys, whitespace, home key, etc., we would have determined more conclusively the tasks performed. For example, was the user typing a document, or programming; playing a game that required the use of arrow keys, or browsing a webpage and pressing page down to browse? The difference between tasks could assist in explaining the differences in rate and quantity between users. For example, space and enter key laden tasks can be more time consuming from a typing perspective, as the space and enter keys require more time to type (300 ms and 550 ms on average respectively) (Kinkead, 1975; Ostry, 1983). Alternatively, users who are browsing may home more frequently.

Furthermore, anonymization made it impossible to determine error rate, another interesting result that could help design systems. Logging editing keys may have provided context that would indicate when errors occurred, through a series of delete or backspace keys following a number of character and space keystrokes.

Further Studies

Our analyses performed led to further questions concerning how users interact with computers. Further research that would yield valuable insight includes studies of the use of shortcut key chords and non-anonymized or less anonymized keystrokes. Having associated details with semi-non-anonymized keystroke data could eliminate potential artifacts of simple document browsing, web browsing, or game playing and provide novel insights into how people multitask.

As wireless devices become increasingly prolific, research centered on touchscreen input and touchscreen typing would also be fascinating. The difference in these devices between manufacturers and the lack of a quality keystroke logger would constitute a limiting factor for completing this research, however.

Acknowledgments

Portions of this work were funded by ONR (N00014-15-1-2275 & N00014-11-1-0275). Chris Dancy, P. Greg Plumb, and Jon Morgan helped gather this data. We thank our colleagues who gave us their time and keystroke logs, and Ysabelle Coutu and Dan Guzek for comments on this paper.

References

Card, S. K., Moran, T., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.

Card, S. K., & Moran, T. P. (1986). User Technology: From pointing to pondering. Proceedings of HPW '86, Proceedings of the ACM Conference on The history of personal workstations, 183-198.

Card, S. K., Moran, T. P., & Newell, A. (1980). The Keystroke-Level Model for user performance time with interactive systems. *Communications of the ACM, 23*(7), 396-410.

Kieras, D. E. (1988). Towards a practical GOMS model methodology for user interface design. In M. Helander (Ed.), *Handbook of Human-Computer Interaction* (pp. 135–158). Amsterdam: North–Holland Elsevier.

Kinkead, R. (1975). Typing speed, keying rates, and optimal keyboard layouts. In Proceedings of the Human Factors Society 19th Annual Meeting 159-161. Human Factors Society: Santa Monica, CA.

Kukreja, U., Stevenson, W. E., & Ritter, F. E. (2006). RUI—Recording User Input from interfaces under Windows and Mac OS X. *Behavior Research Methods, 38*(4), 656–659.

MacKenzie, I. S. (2002). KSPC (Keystrokes per Character) as a characteristic of text entry techniques. *Proceedings of the Fourth International Symposium on Human Computer Interaction with Mobile Devices*, 195-210.

MacKenzie, I. S., Sellen, A., & Buxton, W. (1991). A Comparison of input devices in elemental pointing and dragging tasks. *Applied Ergonomics, 23*(6), 161-166.

Morgan, J. H., Cheng, C.-Y., Pike, C., & Ritter, F. E. (2013). A design, tests, and considerations for improving keystroke and mouse loggers. *Interacting with Computers, 25*(3), 242-258.

Ostry, D. J. (1983). Determinants of interkey times in typing. *Cognitive aspects of Skilled typewriting*, 225-246.

Salthouse, T. A. (1984). Effects of age and skill in typing. *Journal of Experimental Psychology: General, 113*(3), 345-371.

Salvucci, D. D., & Taatgen, N. A. (2011). *The multitasking mind*. New York: Oxford.

Spink, A., Ozmutlu, H. C., & Ozmutlu, S. (2002). Multitasking information seeking and searching processes. *Journal of the American Society for Information Sciences and Technology, 53*(8), 639-652.

Whisenand, T. G., & Emurian, H. H. (1999). Analysis of cursor movements with a mouse. *Computers in Human Behavior, 15*(1), 85-103.