

Are You Your Own Worst Enemy? Self- Interruption on the Computer

Undergraduate Thesis Extended Abstract

By Jing Jin

Advisor: Laura A. Dabbish

Carnegie Mellon University

School of Computer Science

March 21, 2008

Introduction

Although much research has been conducted on the disruptive effects of external interruptions during task execution, very few have studied interruptions occurring without active outside influence, or self-interruption. Self-interruption occurs when a person switches to a different task or activity, and away from their primary task, without active influence from the environment. For example, a person wants to look up some information for work, he or she switches to the web browser, is intrigued by a link on the open web page, and self-interrupts by navigating to that off-topic link.

Self-interruptions can have several negative consequences for the interrupted primary task. In the example above, the user may forget what the original task was (Mark, Gonzalez, & Harris, 2005). Several negative consequences of interruptions in general include wasting valuable time and causing stress when the person realizes how much time has been wasted (Jett & George, 2003), requiring additional attention and time to return to the primary task's context (Czerwinski, Horvitz, & Wilhite, 2004), and increased susceptibility to repeating already-completed tasks because of a loss of context (Mark, Gonzales, Harris, 2005). In addition, self-interrupted tasks are less likely to be resumed compared to externally interrupted tasks (Mark, Gonzalez, & Harris, 2005). In order to decrease inefficiency caused by these interruptions, we must understand the causes and effects of self-interruption and how the environment can influence the frequency and length of interruptions.

The goal of this project is to research how and why people interrupt themselves when working on a computer, and in doing so, build our understanding of the possible negative and positive consequences of self-interruption. The results will inform future design tools to enhance productivity.

Overview

A complementary set of data collection techniques was employed to analyze the patterns and effects of self-interruption including: observations, retrospective interviews, and activity logging studies. An initial set of exploratory observations were conducted to gain a general understanding of the nature and frequency of self-interruption in a typical working environment. Retrospectives were conducted to understand the nature of online self-interruptions occurring within a web browser, by asking users to recall where they went online, and why they went there, using the browser's history as a guide. To collect detailed, accurate information without bias, an activity-tracking study was carried out with participants' computer activities automatically logged for several weeks.

Using the data from these three kinds of user studies, we have defined five categories of self-interruptions, extending those previously defined for interruptions in general (Jett & George, 2003). This categorization includes defining characteristics, pros, and cons of each type of self-interruption. In addition to the self-interruption framework above, we define a model that distinguishes between interruptions and regular task switching activity. Machine learning algorithms will be applied to the activity-tracking data, using Weka, to define models that distinguish between external interruption, self-interruption, and non-interrupting action.

Observations & Retrospectives

Methods

Observations

Observations were conducted with four participants, all of whom use the computer during the regular workday. Participants were recruited through references and personal contacts. Three of the participants were female, and worked as staff or researcher in an Academic institution. All participants use the computer for more than five hours per day on a normal workday.

Participants were studied individually, and each study lasted about two hours. During the first hour, the participant was asked to conduct his or her business as usual, while the researcher observed and took notes on the tasks they performed. A task is defined as a continuous use of the same device, or a continuous conversation with the same person (Gonzalez & Mark, 2004), revolving around the same information. For the second hour, the participants were asked a set of questions about his or her motivation for performing each task to determine whether the task switch was an interruption, and whether the interruption was external.

The observation technique allowed the researcher to examine the participant's activities in great detail as they occurred. However, the limited length of each observation (one hour) meant that the activities observed were not necessarily representative. In addition, the participant's behavior may have been biased by the observer's presence. To counter some of these limitations, we conducted a set of retrospectives based on web browser history, explained in the next section. The retrospectives allowed us to obtain a sample of behavior over a longer period of time (2 days) based on records of naturally occurring user activity (the browser history).

Retrospectives

Retrospectives were conducted with seven participants, three males and four females, the majority of whom were full- or part-time university students. Each participant was asked to go through his or her browser's history for the prior two days and retrace what tasks were performed and reasons they were performed using a think-aloud protocol method. Each of the participants' actions was categorized into an external interruption, a self-interruption, or a non-interruption. Self-interruptions were further categorized according to their intents, triggers and effects.

The retrospective technique allowed collection of data based on the user's naturally occurring behavior inside the web browser. Prompted by their browser history, users were asked to recall the motivation and nature of their activities in the past. However, this presented the limitation of possibly inaccurate recall because the results obtained were strongly dependent on the accuracy of the user's memory. In addition, the data only contained information about web browsing, which is only one of many possible self-interruption activities. Thus we complemented the observations and retrospective data collections with a more in depth follow-up activity tracking study, which will be discussed in a later section.

Results

From the observations, participants spent an average of 8% of their time on interrupting tasks. Also, the ratio of self-interruption to external interruption was low in the observations, at 0.67, compared to a ratio of 0.90 from the retrospectives. Thus, the low rate of interruption in observations could have been caused by hyper-awareness of the researcher's presence. Even though the frequency of self-

interruptions on the observations may be inaccurate, the causes and effects of those interruptions still provide valuable insight.

From the data obtained by the observations and retrospectives, each self-interruption was examined according to its intention, trigger, duration, and effect. Using these characteristics, five types of self-interruptions were defined: Break, Distraction, Reminder, Obstacle, and Deviation. The categories Break and Distraction were extended from categories with the same name from Jett & George's research on interruptions (Jett & George, 2003).

Break

Breaks are defined as spontaneous recesses from work that interrupts the primary task's continuity. Breaks can alleviate stress and decrease chances of distraction caused by boredom. However, breaks can lead to procrastination, and significant time is needed to refocus to the primary task's context. 30.8% of self-interruptions from the observations and retrospectives were Breaks.

Distraction

Distractions are reactions to minimal external stimuli even though immediate action is not necessary. Distractions help increase mental stimulation, and the enjoyment of a task, while performing routine tasks that do not require significant cognitive processing. Conversely, Distractions cause mediocre behavior on tasks that are mentally demanding, and extra time is needed to refocus to the primary task's context. 23.1% of self-interruptions from the observations and retrospectives were Distractions.

Reminder

Reminders are intentional actions taken on a different task to prevent future negligence of that task. Reminders create a sense of accomplishment by preventing forgetfulness through completion of short tasks. However, Reminders are likely to trigger Deviations, and additional time is still needed to refocus to the primary task's context. 23.1% of self-interruptions from the observations and retrospectives were Reminders.

Obstacle

Obstacles are task switches caused by the need to remove something in the environment that has constantly, or will possibly, distract from the task. Obstacles help increase future productivity by removing possible sources of distraction and allowing better focus on the task. However, a failed attempt to remove a distraction can increase stress. 15.4% of self-interruptions from the observations and retrospectives were Obstacles.

Deviation

Deviations are gradual shifts away from the primary task, which may span over several task switches. Deviations provide implicit breaks to relieve stress, but risk the possibility of forgetting the primary task. Deviations also require significant time to refocus back to the primary task's context. 7.7% of self-interruptions from the observations and retrospectives were Deviations.

Activity Tracking

Methods

The activity-tracking study is being conducted with five participants, all of whom use the computer for work on a daily basis. Participants were asked to install and use a program called RescueTime that automatically logs which programs are being used. For each usage instance of a program, the program name, window title, and start and end time are recorded, along with the URL of the website it is accessing, if any. Participants were asked to run RescueTime on their computer for three weeks. At the end of each week, participants have been asked to fill out a survey about their feelings of productivity for the past week. This survey data will be used to identify and separate any atypical periods of activity, as well as provide knowledge of how much the interruptions were influenced by situation and emotion. This survey data will complement activity logs for the week, providing information about the user's context, workload, and well-being during the previous week of activity.

At the end of the three-week study, participants will be presented with the programs they used most often, and will be asked to tag those programs according to what they are usually used for (e.g. work, music, game). Participants will then be presented with a timeline of their computer usage for the past week, displaying their activity in terms of programs used, along with the tags that were created at the beginning of the session. Participants will be asked to recall what they were doing on the computer and why they were doing it. According to their recollections, the activities will be categorized into external interruptions, self-interruptions, and non-interruptions, similar to what was done in the Retrospectives.

Results

The activity-tracking data will be used to make probabilistic predictions towards whether an event was an external interruption, self-interruption, or non-interruption.

Raw data files will first be processed according to several heuristics to decrease the amount of data and increase its meaningfulness. Switches that last for one second or less will be deleted, as the user probably did not have enough time to mentally process and become aware of the switch before returning to the primary task. If repeated switches occur from and back to a single program, then the program will be marked as being continuously used, while programs used in between the switches are possible interruptions. After these heuristics have been applied, the data will be fed into the timeline for the participant to review.

After the completion of the user activity-tracking sessions, each entry in the data will include program name, duration, programs used preceding it, programs used during it, programs used after it, the type of program, and initial probability of being an external interruption, internal interruption, and non-interruption. Each program will be categorized as a browser, chat, e-mail, editor, viewer, or other program.

The initial probability of being an interruption will be calculated based on several rules:

1. Activities on the same program is more likely to be the same task, thus less likely to be an interruption
2. Identical window titles are more likely to be for the same task, thus less likely to be an interruption

3. For browsers, similar URLs or web page titles are less likely to be interruptions
4. Switching back to a program identical to that used in the recent past is more likely to be an internal interruption or a non-interruption
5. Switching to a chat program is more likely to be an external interruption
6. A long sequence of switching back and forth from two programs within a short period of time is more likely to be an external interruption or a non-interruption
7. Switching to and from editors is more likely to be a self-interruption or a non-interruption

After initial probabilities have been calculated, the data table will be loaded into Weka to calculate the probability of each task switch being an external-, self-, non-interruption, or unknown if the confidence is low.

Conclusion

To understand the causes and effects of self-interruption, observations, retrospectives, and activity tracking user studies were conducted to gather daily computer usage data.

Using the data from these studies, five categories were defined for different types of self-interruption, extending those previously defined for interruptions in general (Jett & George, 2003). This categorization defines characteristics, pros, and cons of each type of self-interruption, which will be useful to the design of systems that mitigate the negative effects of self-interruption.

In addition to the self-interruption framework above, we define a model that distinguishes between interruptions and regular task switching activity, with an accuracy of XX% (TBD). This model highlights some of the characteristics distinguishing self-interruption in our dataset, contributing to our understanding of this phenomenon. In addition, this model can provide the intelligence behind a system that would provide activity centric feedback to a user.

References

- Czerwinski, M., Horvitz, E., & Wilhite, S. (2004). A Diary Study of Task Switching and Interruptions. *Proceedings of CHI 2004* , 175-182.
- Gonzalez, V. M., & Mark, G. (2004). Constant Multitasking Crazy: Managing Multiple Working Spheres. *Proceedings of CHI 2004* , 113-120.
- Jett, Q. R., & George, J. M. (2003). Work Interrupted: A Closer Look at the Role of Interruption in Organizational Life. *Academy of Management Review* , 494-507.
- Mark, G., Gonzalez, V. M., & Harris, J. (2005). No Task Left Behind? Examining the Nature of Fragmented Work. *Proceedings of the ACM Conference on Human Factors in Computer Systems* , 321-330.