Using Eye Gaze Patterns to Identify User Tasks

Shamsi T. Iqbal and Brian P. Bailey
Department of Computer Science, University of Illinois
Urbana, IL 61801, USA
{siqbal, bpbailey}@uiuc.edu

ABSTRACT
Users of today's desktop interface often suffer from interruption overload. Our research seeks to develop an attention manager that mitigates the disruptive effects of interruptions by identifying moments of low mental workload in a user's task sequence. To develop such a system, however, we need effective mechanisms to identify user tasks in real-time. In this paper, we show how eye gaze patterns may be used to identify user tasks. We also show that gaze patterns can indicate usability issues of an interface as well as the mental workload that the interface induces on a user. Our results can help inform the design of an attention manager and may lead to new methods to evaluate user interfaces.

Keywords
Attention Management, Gaze, Interruption Overload, Proactive Computing, Task Modeling.

1. INTRODUCTION
Users of today's desktop interface often suffer from *interruption overload*. Notifications from applications executing in the background of user attention such as browsing assistants [14], instant messaging, and email agents [13], all contribute to this burgeoning epidemic of interruption that pervades the desktop interface. With the emergence of proactive computing, the problem of interruption overload will only become more acute.

Research shows that interruptions disrupt users' task performance [2, 3] and emotional state [1, 2, 19], but that these effects can be mitigated if the interruptions occur during periods of low mental workload [2, 3, 4]. Thus, our long-term research seeks to develop an attention manager that can identify moments of low mental workload in a user's task sequence for interruptions to occur.

To begin developing such a system, we must at least be able to identify user tasks and measure the mental workload induced by those tasks. While our previous work has addressed the issue of measuring mental workload through the use of pupil size [8], this work investigates the use of eye movement data to automatically identify user tasks. By coupling a predictive model of task behavior learned from observing identified tasks with a measure of mental workload, we can not only detect mental workload, but we can also *predict* mental workload in the interface.

Current approaches to identifying user tasks include manual specification [12] and making probabilistic inferences from the user event stream [6, 7, 20]. In this paper, we provide an alternative and complementary approach to identifying user tasks based on analyzing patterns of eye movements. We analyze data collected from twelve computer users performing different computer-based tasks and show that each task has a unique signature of eye movement. Thus, given a training set of eye movements for a set of tasks, a system should be able to classify tasks in real-time by comparing the current signature of eye movements to the training set. In addition, we discuss how patterns of eye movements can be used to help identify usability issues with an interface, e.g., by highlighting moments in a user's task sequence where she cannot locate needed information or the desired control. We also show that gaze patterns can indicate the mental workload that an interface is inducing on a user. Our results can help inform the design of an attention manager and may lead to new methods to evaluate user interfaces.

2. RELATED WORK
In this section, we discuss how our work relates to work in task modeling and to the use of eye movements for analyzing user tasks.

2.1 Task modeling
Horvitz et al. [5] infer a probability distribution over a user's task goals by observing several streams of activity. These streams of activity include patterns of mouse and keyboard actions. However, these events do not show how actively involved the user is in a task – the window focus might be in the task but the user's attention might be elsewhere. Our work shows that the use of eye movement, coupled with information from the event stream, may significantly improve a system's ability to accurately infer a user's task.

In the Lumiere project [6], the authors developed a generalized architecture for an intelligent interface. The authors constructed Bayesian models that sense the user event stream to track users' tasks and their information needs. Our work on eye gaze patterns could provide an additional source of significant information that could enhance the predictive power of these models.

Slaney et al. [17] describe an algorithm to cluster and segment sequences of lower-level user actions into sequences of distinct higher-level user tasks. The algorithm uses text contained in interface windows as evidence of the state of the interaction. This provides information primarily about the interaction in the workspace and not necessarily information about the user herself. Coupled with tracking user events, our work on eye gaze patterns could be used to better classify tasks, better predict the user's mental workload during the interaction, and provide more information about the state of the user.

Zhu and Greiner [20] present a novel method for learning a model for user browsing behavior from a set of annotated web logs - web logs that are augmented with the user's assessment of whether each webpage contains the information required to complete her task. The model is subjective by definition. However, eye movement patterns of the user while browsing can indicate objective information such as usability, ease or difficulty of finding the information which may or may not be reflected in the subjective assessment of the page.

2.2 Eye movement

Just and Carpenter [11] showed that eyes do not wander randomly with well structured and speeded tasks. In our study we show that although people keep their eyes focused on areas that are relevant to successful completion of the task – eyes do not remain focused on areas of interest the entire duration of the task. We also showed that the tendency for eyes to wander decreases with increasing difficulty of the task. Our findings are consistent with [11] where the authors state how eye movements reflect cognitive processes.

Eye movements have been studied to understand processes like reading [10], to infer user intentions [9] and to diagnose medical conditions. Rayner [15] showed that during reading, information is only acquired during fixations when the eyes are still. Our study further shows that for reading tasks, users' eye gaze and fixations indicate the effort the user is imparting for the task.

The eye movement enhanced Translation Support system [18] analyzes eye movements during translation, detects patterns in eye movements and responds appropriately. The Reading assistant [16] uses eye gaze to trigger auditory prompting for remedial reading instruction. The application follows the user's gaze path and highlights the words of the text as the reading proceeds from word to word. The aforementioned systems operate in a limited domain where the interpretation of the eye gaze movements can be reduced to a reasonably sized set. Our work is different in that we are investigating the use of eye gaze patterns to identify tasks that a user performs independent of the underlying application, e.g., reading the content of an email message or reading the content of a web page.

3. METHODOLOGY

In our user study, we aimed to investigate whether different eye movement patterns exist for different tasks, independent of the application the task was associated with. For the study, we strived for tasks that would be distinct in the ways users approach them.

3.1 Task Categories

For our user study, we developed the experimental task categories based on a literature review, an informal questionnaire to several users, our own experience, and the consideration of time for the study. Four task categories were developed, each with two difficulty levels (easy vs. difficult):

- *Reading Comprehension.* A user read a given text and answered questions. The easier task belonged to grade 9 level and the more difficult task belonged to grade 17.
- *Mathematical Reasoning*. A user performed math calculations. For the easier task, a user had to mentally add two four digit numbers and select the correct answer from a list of three options. For the more difficult task, a user had to mentally add 4 five-digit numbers, retain the result in memory, and decide whether the result exceeded a given number.
- *Searching*. A user searched for a product from a list of similar products according to specified constraints. For the easier task, a user had to find the product from a list of seven products according to one constraint, e.g., the cheapest camera. For the difficult task, a user had to identify the product using three constraints: the cheapest 3MP camera with 3X digital zoom.
- *Object Manipulation.* A user had to drag and drop email messages into appropriate folders. The user was given a list of emails, four folders, and classification rules. For the easier task, the rule was simple and specific, such as using the size of the email (1K, 2K, or 3K) in the list. For the more difficult task, the rules were less specific, such as the use of topics (travel, course related, fun and humor, announcements). The user had to make a judgment using the information provided in the email.

Although this set of tasks is certainly not exhaustive, we believed that this set was sufficient for the exploratory nature of our work.

3.2 Subjects and Equipment

Twelve computer users (6 female) volunteered in the user study. The average age of the users was 23.7 years (SD = 3.23). As a user performed tasks, we recorded eye movement data using a head-mounted SR Inc., Eyelink II eyetracker with a 250 HZ sample rate and 0.005 degree spatial resolution.

3.3 Experimental Design

The study was a 4 Task Category (Object manipulation, Reading, Mathematical reasoning and Searching) x 2 Difficulty (Easy and Difficult) repeated measures within-subjects design.

3.4 Procedure

Upon arrival at the lab, the user was set up with the eye-tracker and went through a calibration process. The user had to perform 8 tasks – one easy and one difficult for each of the 4 categories. At the beginning of each task category, the user was presented with specific instructions to that category and a practice trial to become familiar with the task. After completing each task category, the user was asked to rate difficulty on a 1-5 scale (5 = very difficult, and 1= very easy). The presentation order of task category and tasks within each category were randomized. The users were instructed to perform the tasks as quickly and as accurately as possible. The system logged task performance and we video recorded the screen interaction for later analysis.

3.5 Measurements

The software application of the eyetracker collected information on the user's gaze, saccades, fixations as the user performed activities on the experimental computer and saved it in real time in a text file. The user's eye movement information as well as the user's on-screen activities were recorded separately. These two data sets were synchronized based on correlating timestamps.

4. RESULTS

For each trial within each task, we identified the main areas of interest (AOI). Figure 1 shows a sample mark up of the AOI of the Search task. Our analysis included the following:

- We analyzed each user's eye movement data to calculate what percent of the total task time their eyes were on the AOI. We call this 'percentage time spent on AOI'.
- We analyzed the eye movement information for each trial for each user to see what percentage of the time spent on AOI their eyes were fixated on each individual AOI.
- For each trial, we plotted the horizontal and vertical coordinates of the eye gaze positions of each user in order to find patterns in how the users accomplished these tasks.

All coordinates in the computer are measured with the top left corner of the screen being the origin (0, 0).

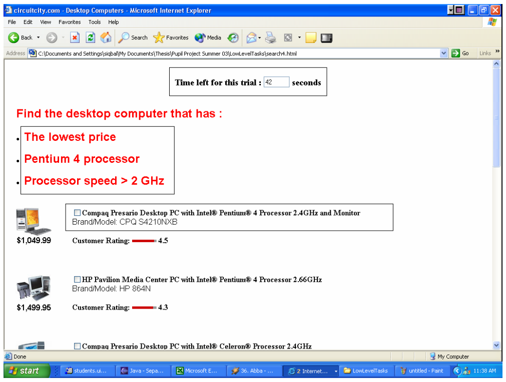

Figure 1. Sample Search task with outlines showing the areas of interest. Outlines were not visible to users.

Tasks	Difficulty	Avg. % Time	S.D.
Reading Comprehension	Easy	50.6	0.13
	Difficult	55.9	0.096
Mathematical Reasoning	Easy	29.4	0.208
	Difficult	35.6	0.269
Search	Easy	48.7	0.089
	Difficult	77.1	0.055
Object Manipulation	Easy	44.8	0.097
	Difficult	61.8	0.108

Table 1. Average percent and standard deviation (s.d.) of time spent on the AOIs averaged across all users.

4.1 Percentage time spent on AOIs

We ran a Repeated Measures ANOVA on the percentage time spent on AOI with task categories and difficulty level as the two factors. The results showed that the users did not keep their eyes on only the AOI throughout the duration of the task. For the easier trials, the users' eyes would move around the screen more than the difficult trials ($F(1,11) = 68.0$, $p <= 0.0001$), which means that there was a significant difference in the eye movement between the easy and difficult trials within each task category.

Also the fixation on AOIs varied across task categories ($F(3,33) = 10.344$, $p <= 0.0001$). Table 1 shows the average percent and standard deviation for time spent on the AOIs across all users.

4.2 Eye gaze fixation within each AOI

From the eye movement data we discovered that users did not view certain parts of some AOIs. Table 2 summarizes the tasks, difficulty levels, and percentage time spent within each AOI. We omit AOIs where the users did not spend significant time, thus some sections in Table 2 do not sum to 100 percent.

For the reading tasks, users spent most time answering the questions. This suggests that users skimmed through the material but had to recall what they read in order to answer the questions. For the easier reading task, users spent 18% of the task time checking how much time was left for the task, while checking time left only 7% of the time for the more difficult reading task.

For the reasoning tasks, users checked the time left more frequently; in the easier tasks, 33% of the task time was spent on checking the time left whereas in the more difficult tasks, 55% of the task time was spent on checking the time left. The main task, i.e., adding the numbers, explained most of the remaining time.

For the search tasks, the majority of the task time was consumed by looking at the objects and their descriptions. The users also spent some time looking at the search criteria. For the easier tasks, the users did not spend more than 16% of the total time looking at the criteria. For the more difficult tasks, the users spent 35% of the task time looking at the criteria which was more complex.

For the easier object manipulation tasks, users spent 34% of the task time checking the criteria for placing the object in the appropriate folder. For the more difficult tasks, this percentage was 56%. The users did not necessarily spend more time contemplating the destination folder, however, once the users had identified the appropriate folder, the movement from the source to

Tasks	Difficulty	AOI	% of time
Reading comprehension	Easy	Time left	18.7
		Read text	8.7
		Questions	9.8
		Answers	15.9
	Difficult	Time left	6.9
		Read text	4.7
		Questions	3.9
		Answers	14.7
Mathematical reasoning	Easy	Time left	33.5
		Add	29.4
		Question	28
	Difficult	Time left	54.7
		Add	31.1
		Question	36.8
Search	Easy	Criteria	15.5
		Search area	84.5
	Difficult	Criteria	35.4
		Search area	64.6
Object manipulation	Easy	Time left	3.0
		criteria	33.6
		email size	55
	Difficult	Time left	0.6
		criteria	56.2
		subject	34.9

Table 2. Percentage time spent within each AOI

the destination followed naturally and the users would view the next object while they were still moving the previous object.

4.3 Eye movements patterns of different tasks

We plotted the horizontal and vertical coordinates of the eye movement for each task across users. As predicted, for each task category, the eye movements were similar across all the users. We show the plots of the data of a single user for simplicity.

The plot for the horizontal coordinate (figure 3(a)) of the eye gaze positions at different time points during the reading task shows a pattern that indicates periodic eye movement from left to right and back to left again. The plot for the vertical-coordinates (figure 3(b)) for the reading task shows a rise in the vertical axes over time, indicating that the eye gaze is moving from top to bottom. Towards the end there is a sudden shift towards a low y-value and then a similar pattern starts. Here the user starts answering questions on a different page and thus reads from the top.

The plot for the horizontal coordinates (see figure 3(c)) for the reasoning task shows that eye movements are restricted to only a few pixels on the screen. This is due to the nature of the task, where the user looks at only a small part of the screen to perform the mathematical calculation. The plot also shows that the user's eyes remain at the same horizontal position for a stretch of time. These are probably the moments where the user is mentally adding the numbers and fixating her eyes at some position on the screen in order to concentrate. As shown by [6], with tasks that are well structured and speeded, people look at what they are working on; the eyes do not wander randomly. This is validated by our data. The plots for the vertical coordinates show (figure 3(d)) a periodic rise and fall. This indicates that the user's eyes are moving from top to bottom as the user adds each column.

The plot for the horizontal coordinates for the search task (figure 4(a)) shows that the user's eyes are restricted to the left side of the screen most of the time. However, there are sudden peaks. Analyzing the videos, we found that this corresponds to the user skimming through the images and prices and if there was a match, then checking for the remaining criteria for a complete match. The plot for the vertical coordinates for the search task (figure 4(b)) shows an increase in the y-dimension as the user moves her eyes down the list. The plot is periodic, which can be explained by the fact that the user has to scroll and therefore the eye moves towards the top part of the screen and moves down again.

The plot for the horizontal coordinates for the object manipulation task show a periodic increase and decrease in the eye coordinates, clearly showing that the user is reading the relevant information left to right (figure 2), moving it to the folder and then moving on to the next object. (figure 4(c)).

The plot for the vertical coordinates for the object manipulation task (figure 4(d)) has an interesting shape – it tapers off slightly as time increases. This is indicative of the task design. As each file is dropped in its folder, everything below it moves upwards and the user has to move her eye less and less vertically. Since the top left corner is the origin, this means that the user's eye is moving towards lesser values of y with time. This is reflected in the plot where the lower bound of the y value remains more or less constant, but the upper value gradually decreases with time.

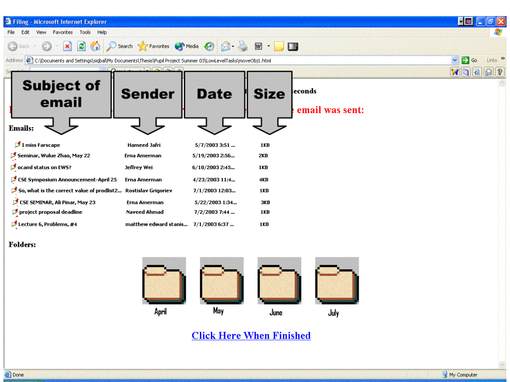

Figure 2. Position of different types of information in the object manipulation task

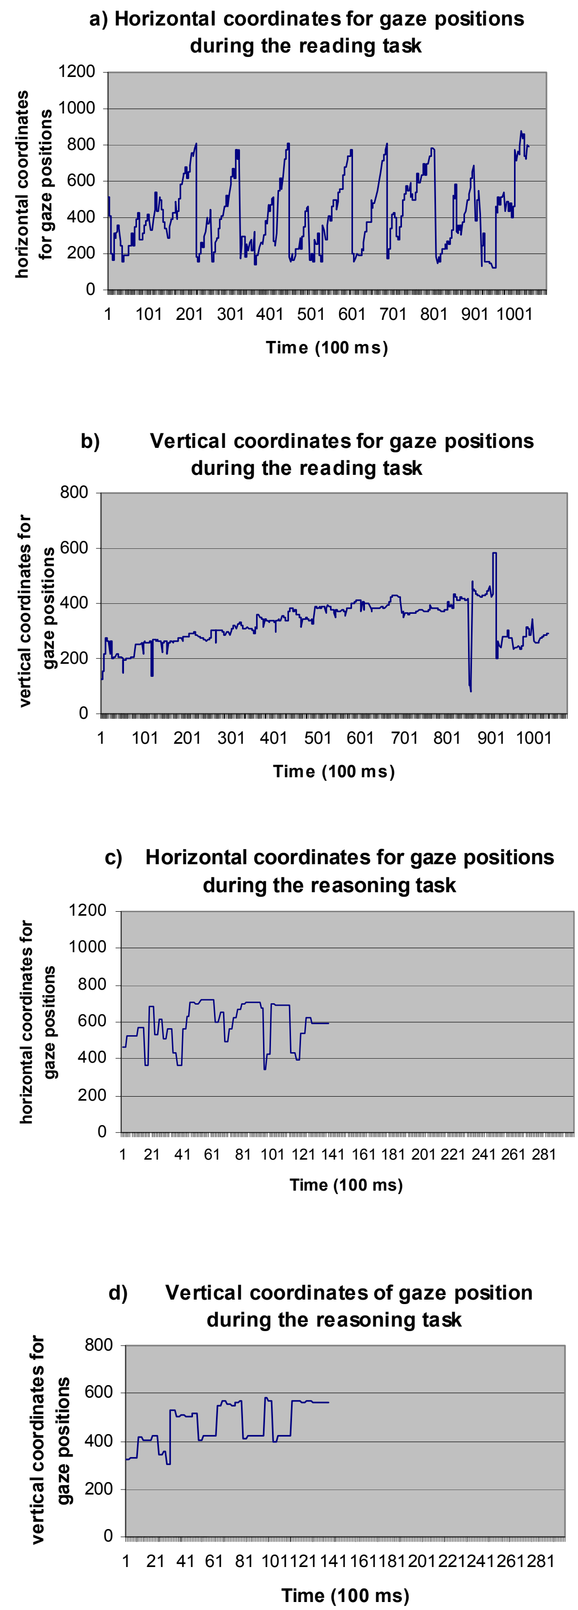

Figure 3(a)-(d). Sample plots of coordinates of the eye movement for the reading and the reasoning task

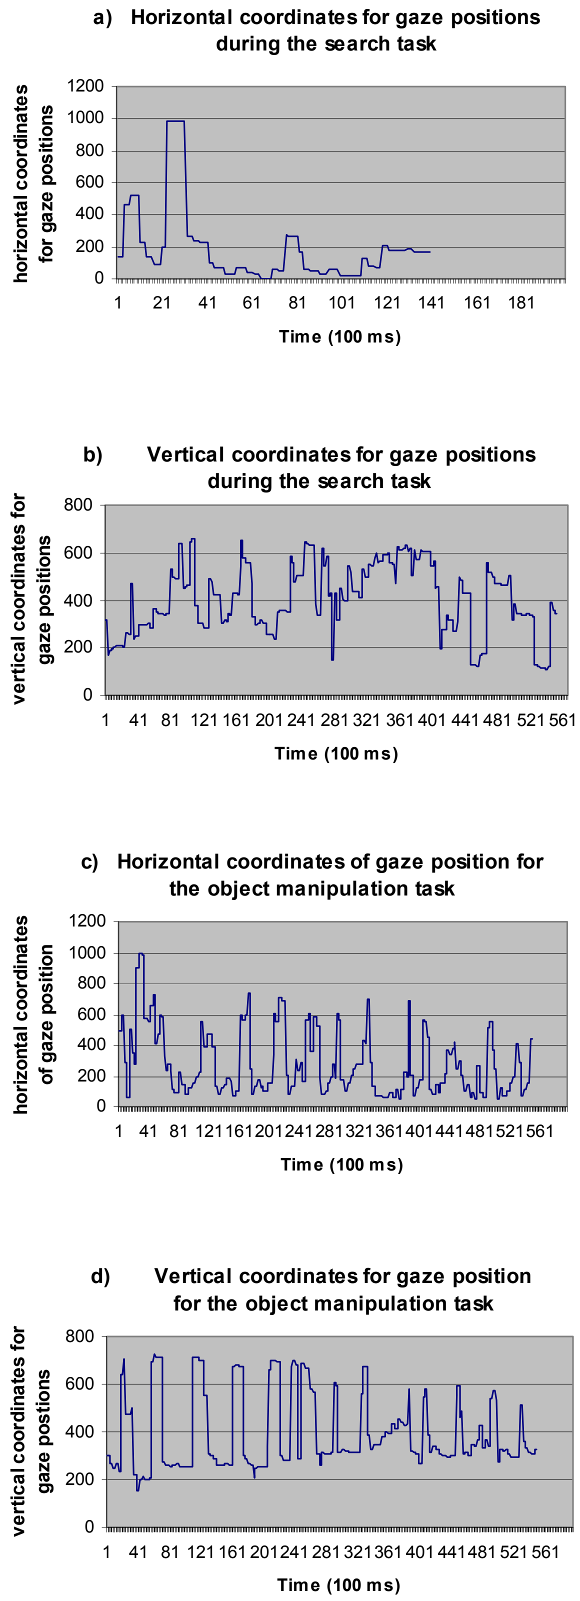

Figure 4(a)-(d). Sample plots of coordinates of the eye movement for the searching and the object manipulation tasks

5. LESSONS LEARNED

From our study, we learned that:

- *Eye gaze patterns can be used to classify the type of task the user is performing.* Our results showed that each task has a unique signature of eye movement. Thus, if there is prior knowledge about the categories of tasks a user typically performs and training sets of eye movement data for those tasks are available, then we can develop a system for classifying user tasks in real time. Classifying user tasks in real time is necessary for making informed decisions about when to interrupt a user engaged in a primary task.
- *Eye gaze patterns are affected by the difficulty of the task.* Our study showed that a user's eyes did not always remain focused on the task. Often the eyes would move around parts of the screen irrelevant to the task. However, this was more of an issue for easy tasks, where a user would look at the areas of interest less than 50% of the task time on average. For the difficult tasks a user would look at the screen 70% of the task time. This suggests that for more difficult tasks, users tend to concentrate more on specific areas of the screen compared to easier tasks. This observation has implications for usability issues of the interface. Patterns of eye movements can be used to highlight moments in a user's task sequence where she has problems understanding the needed information or the desired control.
- *The amount of time a user's eyes are fixated on each area of interest indicates the complexity of that area.* The results showed that a user spends more time on AOIs that induce higher mental workload. This observation is consistent with the observations made in [10]. The amount of time a user's eyes are fixated on any AOI can provide a metric on the complexity of that area and how much mental workload it is inducing. This can be used to evaluate the cognitive complexity of an interface.
- *Eye gaze can determine where and how to present notifications.* Since eye gaze data indicates where a user is focusing her attention and also how complex the task is, it can determine where to display the interrupt notification with minimal disruption to the user. Also, observations of where a user allocates visual attention while not performing any particular task may provide effective locations to display notifications.
- *Users do not always look at all areas of interest.* Our results showed that users occasionally ignored certain important areas of the interface. This means either those parts are not necessary for successful completion of the task or there must be more visual emphasis on those parts of the task to draw a user's attention to them. This observation has implications for evaluating the usability of interfaces. Eye gaze data can be used to evaluate user interfaces to see whether all functionalities of the interface are visible to the user or whether they are used often enough to validate their placement in the interface.
- *Eye gaze can guide the visual organization of a task.* Based on our results, we conclude that eye gaze provides a valuable source of information on the difficulty of the task and what areas of the screen the user is focusing her attention on. This can help guide the design of the visual layout of a task. Controls that are related may be grouped together for better performance, so that the user does not have to search too far to locate desired controls or required information.

The lessons learned from this study have important implications for the design of an attention manager as well as user interface design. An attention manager attempts to balance a user's need for minimal disruption with an application's need to effectively deliver information. Identifying in real time what task the user is performing is essential for an attention manager to detect and predict user tasks. Our findings show eye gaze patterns to be a promising source of information about a user's current task. This information can be coupled with system events to provide a more accurate identification of the user's current task. Finding the appropriate location to display notifications to the user is another relevant issue for the attention manager. Eye gaze direction shows where the user's visual attention currently is, indicating a good location for notification display. Eye gaze patterns can help identify usability issues in an interface by showing where in an interaction sequence a user struggles to locate needed information. Eye gaze patterns can also suggest points in an interaction sequence that induce too much cognitive complexity on the user.

6. FUTURE WORK

We plan on developing computational methods for comparing short periods of eye movement data to a training set of defined tasks. We plan to convert the time domain data to the frequency domain using a Fourier transform and then compare the resulting coefficients. We also want to investigate including mouse and keyboard activity to further improve the classification. Pattern analysis techniques such as Principal Component Analysis may also help us identify distinct patterns in different dimensions.

Our findings in this study will help us develop a better task model for predicting a user's tasks, which is an important component of our attention manager. When coupled with a measure of mental workload, the model of user task behavior learned over a period of time can be trained to predict moments of low mental workload in a user's task sequence. The attention manager can defer interruptions from applications until the next opportune moment in the task sequence. We believe that an effective attention manager can increase user productivity and decrease frustration, annoyance, and anxiety - enhancing the overall interaction experience for users of the desktop interface.

7. REFERENCES

1. Adamzcyk, P.D. and B.P. Bailey. If Not Now When? The Effects of Interruptions at Various Moments within Task Execution. *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2004, to appear.
2. Bailey, B.P., J.A. Konstan and J.V. Carlis. The Effects of Interruptions on Task Performance, Annoyance, and Anxiety in the User Interface. *Proceedings of Interact*, 2001, 593-601.
3. Cutrell, E., M. Czerwinski and E. Horvitz. Notification, Disruption and Memory: Effects of Messaging Interruptions on Memory and Performance. *Proceedings of Interact*, Tokyo, Japan, 2001, 263-269.
4. Gillie, T. and D. Broadbent. What Makes Interruptions Disruptive? A Study of Length, Similarity, and Complexity. *Psychological Research*, 72(1989): 243-250.
5. Horvitz, E., A. Jacobs, and D. Hovel. Attention Sensitive Alerting, *Proceedings of the Conference of Uncertainty and Artificial Intelligence*, 1998, 305-313.
6. Horvitz, E., J. Breese, D. Heckerman, D. Hovel, K. Rommelse. The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Madison, Wisconsin, 1998, 256-265.
7. Huber, M.J. and Simpson, R. Plan Recognition to Aid The Visually Impaired. *Proceedings of User Modeling,* 2003, 138-142.
8. Iqbal, S.T., X.S. Zheng and B.P. Bailey. Task Evoked Pupillary Response to Mental Workload in Human-Computer Interaction. *Proceedings of the ACM Conference on Human Factors in Computing Systems,* 2004, to appear.
9. Jacob, Robert J.K. The Use of Eye Movements in Human-Computer Interaction Techniques. *ACM Transactions on Information Systems* 9 (3): 152-169, 1991.
10. Just, M.A. and P.A. Carpenter. A Theory of Reading: From Eye Fixations to Comprehension, *Psychological Review* 87 (4): 329-354, 1980.
11. Just, M.A. and P.A. Carpenter. Eye fixations and cognitive processes, *Cognitive Psychology* 8 (1976): 441-480.
12. Kieras, D.E., S.D. Wood, K. Abotel and A. Hornof. Glean: A Computer-Based Tool for Rapid Goms Model Usability Evaluation of User Interface Designs. *UIST*, 1995, 91-100.
13. King, W.J. Reflections of Thought: Cognitive Facial Expressions in the Human Interface, *Proceedings of the 5th IEEE International Workshop on Robot and Human Communication*, 1996, 195-200.
14. Lieberman, H. Letizia: An Agent that Assists Web Browsing. *IJCAI 95,* AAAI Press, 1995.
15. Rayner, K. Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin* 124 (3): 372-422, 1998.
16. Sibert J.L., Gokturk, M., and Lavine, R.A. The Reading Assistant: Eye Gaze Triggered Auditory Prompting for Reading Remediation, *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2000, 101-107.
17. Slaney, M., J. Subrahmonia and P. Maglio. Modeling Multitasking Users. *Proceedings of User Modeling,* 2003, 188-197.
18. Takagi, H. Development of an eye-movement enhanced translation support system. In *Proc. APCHI'98* (Asian-Pacific Computer Human Interaction Conference), 114-119.
19. Zijlstra, F.R.H., R.A. Roe, A.B. Leonora and I. Krediet. Temporal Factors in Mental Work: Effects of Interrupted Activities. *Journal of Occupational and Organizational Psychology*, 72 (1999): 163-185.
20. Zhu, T., R. Greiner and G. Haeubl. Learning a Model of a Web User's Interests. *Proceedings of User Modeling*, 2003, 65-75.