AROMA: Ambient awaReness through Olfaction in a Messaging Application

Does Olfactory Notification Make 'Scents'?

Adam Bodnar
Dept. of Computer Science
University of British Columbia
2366 Main Mall
Vancouver, BC V6T 1Z4
bodnar@cs.ubc.ca

Richard Corbett
Dept. of Computer Science
University of British Columbia
2366 Main Mall
Vancouver, BC V6T 1Z4
rcorbett@cs.ubc.ca

Dmitry Nekrasovski
Dept. of Computer Science
University of British Columbia
2366 Main Mall
Vancouver, BC V6T 1Z4
dmitry@cs.ubc.ca

ABSTRACT

This work explores the properties of different output modalities as notification mechanisms in the context of messaging. In particular, the olfactory (smell) modality is introduced as a potential alternative to visual and auditory modalities for providing messaging notifications. An experiment was performed to compare these modalities as secondary display mechanisms used to deliver notifications to users working on a cognitively engaging primary task. It was verified that the disruptiveness and effectiveness of notifications varied with the notification modality. The olfactory modality was shown to be less effective in delivering notifications than the other modalities, but produced a less disruptive effect on user engagement in the primary task. Our results serve as a starting point for future research into the use of olfactory notification in messaging systems.

Categories and Subject Descriptors

H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems; H.5.m [**Information Interfaces and Presentation**]: Miscellaneous—*Olfactory Notification*

General Terms

Design, Experimentation, Measurement

Keywords

notification systems, olfactory display, multi-modal interfaces, ambient awareness, user study, HCI

1. INTRODUCTION

Advances in multi-modal interfaces have enabled the development of systems that aim to increase the efficiency with which people can perform multiple tasks. It is common for users of today's multitasking computing environments to desire immediate access to information regarding background events unrelated to their primary tasks. These can range from the receipt of new email, to a change in system status such as low battery power on a laptop, to the online arrival of a friend in an instant messaging system. Paying attention to such events can be useful and important, but it can also adversely affect one's focus on the primary task, inhibiting performance and causing dissatisfaction. Thus, mechanisms used to notify users of these events must effectively draw their attention towards important information while avoiding unnecessary disruption.

Traditional computer interfaces primarily rely on visual notification to convey information to users, with occasional use of auditory signals. For this reason, notification of background events is usually delivered through visual or auditory stimuli. However, this does not take advantage of the fact that humans possess extraordinary sensing capabilities apart from these channels [1]. One such capability is the sense of smell, a powerful and versatile medium that humans use to obtain background information of all kinds, from the danger of a burning fire to the proximity of a loved one. Yet smell remains a poorly understood modality, both in the realm of computer technology in general and in particular as a means of providing background notifications.

It is the purpose of the work described in this paper to investigate the use of smell as a medium for notification. In the experiment described in the following sections we aim to gather pertinent information for future design of notification systems. In Section 2 we present the background for this experiment, both in terms of related work and our motivation. In Section 3 we provide a detailed methodology of the experiment. Section 4 reports the results of our investigation. Section 5 discusses how our results can be applied to future notification system design. Section 6 discusses directions for future research and concludes.

ICMI'04, October 13–15, 2004, State College, Pennsylvania, USA.

2. BACKGROUND

In this section we discuss the research directions and empirical results for notification systems and olfactory stimuli. We also explain the related work behind our experimental design and the motivation for our study.

2.1 Notification Systems

Notification systems provide users with information regarding background events, which are often unrelated to their primary tasks. Examples of such systems include email clients, system resource management monitors, and instant messaging applications. The benefits of notification systems are well established in literature, and previous research suggests that these systems can provide significant value to users [10]. However, studies have also revealed that notifications can lead to user frustration and decreases in primary task performance. The need to deliver real-time information about background events effectively to users without causing unwanted distraction has motivated much recent research into notification systems [2, 3, 6, 9]. Most of this research has focused on the visual and auditory modalities, since these are the most readily available output channels found on today's computer systems. Other possible modalities, such as haptic and olfactory, have been largely ignored. However, there is no concrete evidence to support that either visual or auditory display represents the optimal means of notifying users of background events.

Based on a lack of research in this area, we have decided to explore and evaluate the effectiveness of olfactory notification in the domain of messaging applications. To date, research related to messaging has focused primarily on the disruptive effects of messages to users' focus and performance on primary tasks [2, 3]. Another study has attempted to infer the cost of user interruption by messages through Bayesian modeling techniques [6]. With the exception of this study, the factors examined in the literature have included the duration and content of messaging interruptions [2], the nature of the primary task [3], but not the modality of the interruption. The Horvitz study [6] considered visual and auditory messaging notifications alone and in combination, but strictly in the context of their expected interruption costs, rather than actual disruptive effect or user preferences.

2.2 Olfactory Notification

Compared to other senses such as vision, hearing, and touch, the sense of smell has not been well explored or evaluated in multi-modal interface research. There are a number of reasons for this: olfactory communication is limited in bandwidth, hard to categorize, inappropriate for rapidly changing information, and difficult to provide on demand in a controlled manner [8]. But the ability of a scent to move easily from the periphery of a user's attention to the centre and back again makes smell highly suitable for ambient display of information [11]. Recent work taking advantage of this fact has produced novel techniques for olfactory display of information ranging from the state of the stock market to calendar reminders in Microsoft Outlook [7]. Other research has revealed that olfactory stimuli can promote a feeling of task immersion in a user when used in conjunction with auditory and visual stimuli [4]. While these results are both interesting and encouraging, we believe that the full potential of the olfactory modality for providing notification information remains untapped.

With AROMA, we aim to strengthen the empirical literature on notification systems by evaluating the olfactory modality, as compared to the visual and auditory modalities, as a non-disruptive and effective channel of notifying users of background events.

2.3 Motivation

A previous attempt to compare the effects of olfactory notification to visual and auditory notifications has been carried out by Arroyo *et al.* [1]. However, this study focused only on the relative disruptiveness of each modality and therefore did not attach any semantic content to notifications. While important, we believe this approach disregards the information value that alternate modalities such as smell can provide in notification systems. Clearly, visual and auditory notifications can convey informative messages to users, but their ability to cause interruptions is well documented [2]. Yet this may be due to the disruptiveness of these modalities rather than to any inherent requirement that messages force the user's attention to be diverted. Indeed, messaging systems can avoid this by utilizing some kind of ambient display, as described, for instance, in [5]. Since we are not aware of any previous comparisons of the effectiveness of multiple modalities in delivering messaging notifications, our intention is to evaluate the extent to which olfactory notification can effectively support such *ambient awareness* compared to visual and auditory display.

Our theoretical framework for evaluation is based on a model for classifying notification systems outlined in McCrickard *et al.* [9]. This model categorizes notification system user goals and designs according to three criteria: *interruption* (the degree of transfer of focus from a primary task to the notification), *reaction* (the desirability of a rapid and accurate response to the notification), and *comprehension* (the suitability of the notification system to providing information cues that can be remembered and understood at a later time). We postulate that the properties of the olfactory modality are conducive to its use in notification systems, which de-emphasize the interruption dimension in favour of providing reaction and/or comprehension. McCrickard describes three types of such systems requiring varying proportions of these elements: ambient media (low reaction, high comprehension), indicator (low comprehension, high reaction), and secondary display (some reaction and comprehension). *Ambient awareness* of messaging notifications as described above falls into this latter category, and therefore our evaluation focuses on the effectiveness and disruptiveness of smell as a secondary display mechanism in the context of messaging notifications.

A practical situation where a notification system would need to provide both reaction and comprehension would be a scenario where the user needs to make a distinction between important notifications that must be immediately attended to and notifications that the user would prefer to ignore. An example of such a scenario is experienced by graduate students, who often collaborate on documents. A graduate student who is collaborating with another student may want to immediately respond to any instant messages they receive from their partner as the team attempts to meet a quickly approaching deadline. However, the same graduate student may wish to ignore any messages from other friends in order to better focus better on the primary task. We have attempted to explore this distinction in the experiment described in the remainder of this paper.

3. EXPERIMENT

We designed an investigation into differences in user performance and preference as a function of notification mode.

These differences were measured in an experiment that consisted of a primary task for participants to focus on while they received either visual, auditory, or olfactory notifications that could have one of two possible meanings, as in the scenario described in the previous section.

3.1 Goals

Our experiment is designed to answer the following questions:

- Can olfactory notification be used to make a user aware of an event?
- How does the mode of notification affect a user's primary task performance?
- How does the mode of notification interruption affect a user's ability to correctly respond to a notification?
- Which mode of notification do users prefer?

3.2 Task

The primary task that was chosen for the experiment was the completion of arithmetic questions similar to skill-testing questions found on lottery tickets. This task was chosen, rather than the collaborative task of the form described in Section 2.3, because it allowed the authors to get periodic input from users that could easily be deemed as correct or incorrect. By gathering input through the completion of smaller tasks rather than one long task, we could accurately track participants' progress during the experiments. This task proved to be successful in engaging the focus of the participants for the duration of the experiment. Furthermore, the task could not be accurately completed without devoted concentration.

3.3 Notification Conditions

During the experiment, participants could receive one of two types of notifications:

- Type 1 - When participants received this notification, they were instructed to immediately interrupt what they were doing and attend to it. This involved recording the time as it appeared in the experimental application into a text box, as shown in Figure 1. Once this was done they could return to their primary task.
- Type 2 - When participants received this notification, they were instructed to ignore it and continue with the primary task.

The duration of the notifications depended on participant response. If the participant responded to a notification, it would cease immediately after the 'Submit Time' button was clicked regardless of its type. Otherwise, the notification would continue for 10 seconds. The behavior of Type 1 and Type 2 notifications was designed to be the same in order to avoid confusing participants if they responded incorrectly.

3.4 Pilot Study

A pilot study was conducted with four participants in order to evaluate preliminary choices of notifications, as well as the proposed format of the experiment. From this work we improved our design and settled on the final task and stimuli.

Figure 1: A screen shot of the application. The list of arithmetic questions is visible on the left. On the right are the areas for entering responses to questions and entering the current time as a response to Type 1 notifications. The dark rectangle is a visual notification used in the experiment.

Based on the results of the pilot, participants would be asked to complete four blocks of questions each five minutes in length. The first block would be a training phase where participants would not receive any notifications. During each of the other three blocks participants would receive notifications from one of the three modalities. We chose this time constraint to avoid fatigue effects that might be experienced by participants after long periods of answering questions that would be detrimental to their performance.

In the pilot study, we experimented with different frequencies of notification delivery, ranging from two to ten notifications per block of questions. Based on the results of the pilot, we chose to deliver six notifications per block (three of each type) in order to meet the constraints of our olfactory dispensers and to produce a significant effect on participants' performance while allowing them to concentrate on the primary task for the majority of the experiment.

It was determined in the pilot study that some questions of varying length could greatly differ in difficulty, thus we modified the arithmetic questions so that each question consisted of five basic operations. A minimum of 30 questions were used in each block to ensure that no participants ran out of questions to complete within the time allotted.

Based on the results of the pilot study, we selected the following stimuli for each notification modality:

- Olfactory - From ten different scents we selected two scents that could be easily differentiated between at intervals of roughly 30 seconds. One of these was an extract of cloves, while the other was VitalActivePlus, a mixture of essential oils which had an artificial eucalyptus scent.
- Visual - We selected two static solid-coloured rectangles, blue for one type of notification and red for the other.

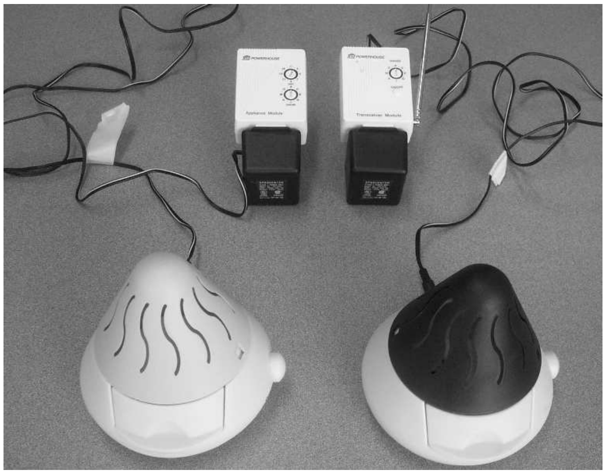

Figure 2: Scent diffusers and X.10 components used to provide olfactory notification. The diffusers were augmented with directional funnels for efficient scent release.

Figure 3: Experimental setup of olfactory diffusers.

- Auditory - We selected two differently toned bell-like sounds. These sounds were chosen because they were easy to differentiate between and terminate when necessary.

3.5 Apparatus

The experiment was conducted on a Fujitsu-Siemens P3-800 laptop with 256 MB of RAM running MS Windows 2000. The experimental system software was fully automated and implemented as an application in Visual Basic 6. Visual notifications were activated within the application, while auditory notifications were played by the application through system speakers.

To deliver the olfactory notifications we employed two Spa Scenter diffusers (shown in Figure 2) and augmented them with directional funnels to provide rapid and controlled scent release (shown in Figure 3). This also helped to limit the amount of scent we had to release to notify the participants, thus making the experiment site less saturated with scents. The diffusers were controlled externally through an X.10 serial interface activated by the application.

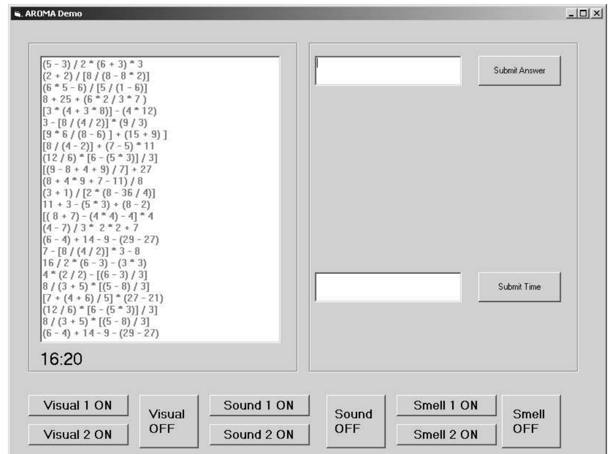

Figure 4: The depicted application was used to familiarize the participants with the notifications before each block.

The X.10 switchable components made audible 'clicks' when they were activated and deactivated. These sounds were successfully mitigated by placing the components within a padded container, as well as requiring participants to wear Maxtor sound blocking headphones throughout the experiment. Furthermore, to be sure that no additional auditory cues were available to the participants, a third scent dispenser, producing no scent, was placed next to the participants and remained activated throughout the experiment. Its purpose was to mask any fan sounds that might be produced by the active diffusers.

Prior to the experiment, participants were introduced to the experimental application interface (shown in Figure 1) and were provided with instructions on its use. Participants were informed that they would be asked to solve a number of arithmetic questions, giving them a task to focus on while they received different notifications. Participants were asked to answer the arithmetic questions in a timely and accurate manner, but were given no other external motivation, such as a monetary reward, to perform well.

After the introduction, participants completed a five minute training session, where they answered arithmetic questions in a manner that mirrored actual experimental conditions, except that no notifications were provided. The results of the training phase were used to provide a basis of comparison to experimental blocks that contained notifications.

3.6 Procedure

After the training session, participants were instructed to complete three blocks of arithmetic questions, one block for each notification mode. Prior to the experiment, a detailed explanation of the interface as well as the two types of notifications was provided. Prior to each block, participants were introduced to the mode of notification that would be provided (through the application shown in Figure 4), and given as much time as they needed to feel comfortable in their ability to differentiate between the two types. Before each block, participants were instructed that they would have five minutes to answer as many arithmetic questions as they could and instructed to complete the questions in the order they appeared on the screen. Participants were also informed that they would be randomly interrupted by a given mode of notification.

Participants were instructed in the use of the sound blocking headphones, and asked to use them for the course of the experiment. The investigators then clarified any questions the participants might have, and explained that once the formal experiment had begun that no further questions would be answered. No investigators were present in the room while participants conducted the experimental task, however a video camera was set up behind the participants to record their progress.

3.7 Measures

We utilized both quantitative and qualitative measures to compare the three notification modes. Our main raw dependent variables for the quantitative measures were number of questions completed and number of errors. In addition, due to the variability in the number of questions completed between subjects, we used two additional dependent variables: the number of questions completed as a fraction of the number completed during the training session, and the error rate (i.e. the number of errors as a fraction of the number of questions completed). We also recorded the participants' responses to notifications and analyzed their correctness and completeness (i.e. whether participants missed any Type 1 notifications or responded to any Type 2 notifications).

A questionnaire and an informal interview were administered after the completion of the experiment. The questionnaire asked participants to rank each condition on a five-point Likert scale with respect to *effectiveness* (defined as ease of responding to Type 1 notifications) and *non-disruptiveness* (defined as ease of ignoring Type 2 notifications), and to rank the conditions by both of these criteria. The interview attempted to elicit the reasons for these preferences and solicit comments and suggestions about the modalities and the experiment.

3.8 Design

The design used was a within-subjects design with three conditions corresponding to notification modalities. Within each condition, subjects experienced a total of three Type 1 notifications and three Type 2 notifications. A within-subjects design was chosen for its increased statistical power and because it allowed for comparative participant comments on the three modality conditions. To minimize training and fatigue effects, conditions were counterbalanced using a Latin square, resulting in three condition configurations. Sets of arithmetic questions were counterbalanced to ensure that the relative difficulty of each set was not a confounding factor. Furthermore, the two notifications for each mode were counterbalanced to eliminate any potential difference in effectiveness that may have occurred as a function of colour, pitch, or smell.

3.9 Participants

Participants were graduate students in Computer Science at the University of British Columbia. A total of 12 participants were involved, four per condition configuration. All were between the ages of 22 and 37; there were 10 males and 2 females. All described themselves as intermediate or expert computer users; all but one described themselves as having some or regular experience with computer applications involving notifications, such as instant messaging. None of the participants took part in pilot studies. All participants reported normal hearing, normal or corrected to normal vision, and normal sense of smell without any temporary impairments such as sinus congestion.

Each participant was involved for approximately 45 minutes. No compensation was offered to participants for their time. Participants were encouraged to perform to the best of their abilities, but no motivation was offered to those who performed better than others.

3.10 Hypotheses

Our hypotheses were as follows:

- **H1**: Olfactory notification results in equal, or more, questions completed than visual and auditory notification.
- **H2**: Olfactory notification results in a lower error rate than visual or auditory notification.
- **H3**: Olfactory notification is at least as effective as visual and auditory notification.
- **H4**: Olfactory notification is less disruptive than visual or auditory notification.

4. RESULTS

In this section, we present the results of our user study and provide a summary of these results based on our hypotheses.

4.1 Performance

To determine whether there were differences in the correctness and completeness of participants' responses to notifications of different modalities, we compared these to records of notifications actually sent. All visual and auditory notifications were responded to correctly, with the exception of one mistakenly ignored Type 1 auditory notification (combined error rates of 0.0% and 1.4%, respectively). Participants mistakenly ignored three Type 1 olfactory notifications (two by the same participant) and incorrectly responded to two Type 2 olfactory notifications, for a combined error rate of 6.9%. A one-way analysis of variance (ANOVA) found no statistically significant difference in notification response error rates between modalities.

To determine whether there were differences in performance between the notification modes, one-way ANOVA's were calculated for each of the dependent variables: number of questions completed (raw and as a percentage of the number completed during the training session) and number of errors (raw and as a fraction of the number of questions completed). The means and standard deviations for each of these dependent variables by mode are shown in Table 1. There were no significant differences in raw number of questions completed, raw number of errors, or error rate between the modes. However, there was a statistically significant difference in the number of questions completed as a fraction of those completed during the training session ($F(2,35)=4.491$, $p=.019$)).

To examine this effect in more detail, we compared the number of questions completed as a percentage of the number completed during the training session for each possible pair of modes using the Bonferroni method. The comparison revealed that the percentage was significantly higher for

Table 1: Means and standard deviations for performance measures (N=12).

Dependent Variable	Auditory		Visual		Olfactory	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Number of questions completed	12.83	3.738	12.25	2.800	10.08	2.937
Number of questions completed as fraction of training phase	1.241	0.193	1.202	0.225	0.988	0.249
Number of errors	2.500	1.732	2.4167	0.900	2.000	1.206
Error rate	0.2005	0.1367	0.2144	0.1115	0.2183	0.1516

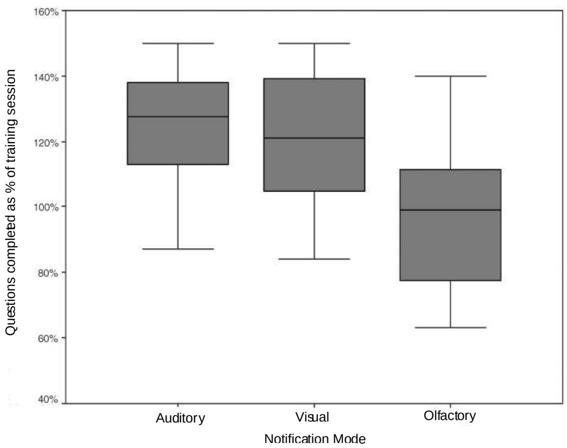

Figure 5: Boxplot of notification type versus performance (N=12). The line through each box represents the median.

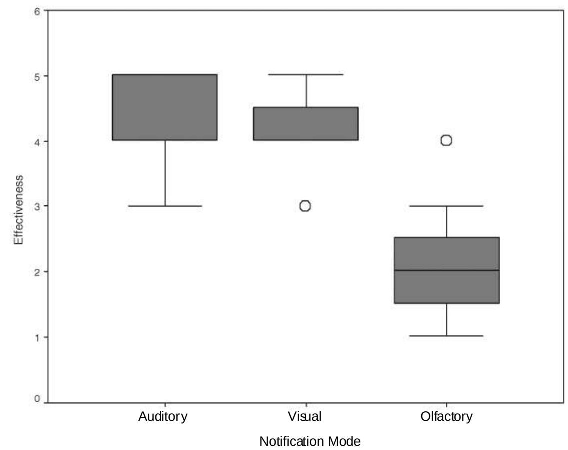

Figure 6: Boxplot of notification type versus effectiveness (N=12). Circles represent outlier responses.

auditory compared to olfactory (mean difference = 25.4%, p=.027). The overall results for the number of questions completed as a percentage of the number completed during the training session are shown in Figure 5. These results suggest that, relative to their performance without notifications, participants were able to complete significantly more questions when responding to visual and auditory notifications than when responding to olfactory notifications.

To determine whether the performance was subject to a training or fatigue effect, one-way ANOVA's were also calculated for each of the dependent variables mentioned above with the order in which blocks appeared rather than notification mode as the independent variable. No statistically significant differences were found, indicating a lack of evidence for learning or fatigue effects.

4.2 Self-Reported Measures

For the two self-reported measures (effectiveness and non-disruptiveness), we analyzed participant rankings of each notification mode both individually and relative to each other.

The results of the individual rankings are summarized in Figures 6 and 7. To determine whether there were differences in effectiveness and non-disruptiveness rankings between the modes, one-way ANOVA's were calculated using each as the dependent variable. Significant differences were found for both effectiveness (F(2,35)=4.112, p=.025) and non-disruptiveness (F(2,35)=34.685, p<.001). Post-hoc Bonferroni comparisons revealed that participants considered olfactory to be significantly less effective than both auditory (mean difference = 2.42, p<.001) and visual (mean difference = 2.08, p<.001), but significantly less disruptive than auditory (mean difference = 1.42, p=.029). Participants also considered olfactory to be less disruptive than visual, but this difference was not statistically significant (mean difference = 1.08, p=.131).

For each of the relative rankings, we analyzed the frequency with which each mode was ranked first. In order to do this, the Chi-square statistic was calculated to determine if the actual frequencies were significantly different from expected equal frequencies. Chi-square was significant for both effectiveness and non-disruptiveness. Nine of the participants perceived auditory to be the most effective condition (χ^2(2,12) = 10.5, p = .005), with the rest choosing visual. Eight of the participants perceived olfactory to be the least disruptive condition (χ^2(2,12) = 6.5, p = .039), with all but one of the rest choosing visual.

Comments that participants brought up during the post experiment interviews reflected a diversity of preferences. Four participants mentioned that the two types of olfactory notifications were difficult to differentiate between and thus required concentration to process. However, others saw their relative subtlety as an advantage in terms of non-disruptiveness. Three participants indicated that their difficulty with distinguishing scents was at least partially due to the lack of a mechanism for quickly removing scents from the experimental site after a notification. Several partic-

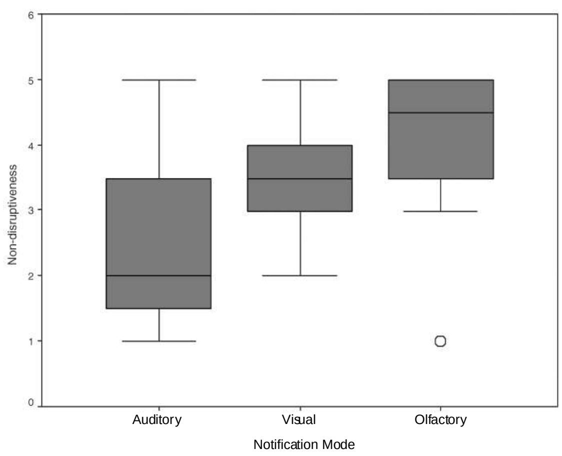

Figure 7: Boxplot of notification type versus non-disruptiveness (N=12). Circles represent outlier responses.

ipants also noted their lack of training and experience in distinguishing scents, and suggested that their performance and satisfaction with scent notification could improve given more practice. This suggestion was confirmed by a participant with previous experience of aromatherapy, who preferred olfactory notification and did not see it as a hindrance in terms of performance on the primary task.

4.3 Summary of Results

A summary of our results according to our hypotheses follows:

- **H1**: Olfactory results in fewer questions completed than visual or auditory.
- **H2**: Olfactory does not result in lower error rate than visual or auditory.
- **H3**: Olfactory is less effective than visual or auditory.
- **H4**: Olfactory is less disruptive than visual or auditory.

5. IMPLICATIONS FOR DESIGN

Our experimental results have significant implications for notification system and multi-modal interface design. We have shown that, for a high concentration task, auditory and visual notifications yield significant performance benefits in terms of increased user productivity on the primary task. It is our belief that this is due to the lack of user familiarity with olfactory notification. All but one of our participants described themselves as being familiar with current messaging technology, which makes heavy use of both auditory and visual notifications. Therefore it is reasonable to believe that our participants had more experience with both auditory and visual notifications than with olfactory notification. We postulate that with proper training and time, olfactory notification would result in similar primary task performance. Thus, a notification system that aims to maximize primary task performance should make use of auditory or visual notifications, unless proper training in the use of olfactory notification can be provided.

Our results have also shown that users find the olfactory modality to be the least effective mode for notification while performing a high concentration task. We did not encounter the difficulties in providing on demand olfactory communication in a controlled manner reported in [8]. Rather, our results suggest that one of the main shortcomings of our system was its inability to remove the dispensed scents in a manner that supported the nature of the task and frequency of the notifications. Thus, a notification system that aims to provide frequent notifications could make use of olfactory feedback, as long as there exists a method for clearing the dispensed scent from the user's immediate vicinity.

Finally, we have shown that olfactory notification is perceived to be significantly less disruptive than auditory notification, and overall to be the least disruptive of the three modalities we investigated. This property could enable users to remain more focused on their primary task and ignore unimportant notifications that would be disruptive if provided through the auditory or visual modalities. This is an encouraging result for future research, since it implies that the olfactory modality would be highly appropriate for use in notification systems that aim to minimize their level of disruptiveness.

6. CONCLUSIONS AND FUTURE WORK

An important challenge for notification systems is to effectively draw users' attention towards important information while avoiding unnecessary disruption. In this paper, we have presented an experiment comparing the effectiveness and disruptiveness of three different modes for delivering notifications: auditory, visual, and olfactory. Our quantitative results indicate that olfactory notifications negatively impact primary task performance, a finding confirmed by participant feedback. However, the majority of participants found olfactory to be the least disruptive of the three modalities as a notification medium. Several indicated that their relatively poor performance was likely caused by a lack of experience with receiving and interpreting scent notifications. Based on these results, we infer that with greater training and experience, olfactory notification could result in similar task performance to visual and auditory notification while remaining the least disruptive modality.

The results presented motivate several future research directions for olfactory notification. We would like to explore what scent properties, such as composition, frequency, and intensity, are optimal for providing olfactory notifications, as well as various means of effectively removing scents after they are diffused. We are also interested in determining how well the results of this work scale to a greater number of notification types and whether they would generalize to primary tasks with other properties (e.g. collaborative writing or video games). In addition, we would like to investigate how much time and training is required to remove the negative impact that olfactory notification has on primary task performance. Finally, our background research uncovered a lack of guidelines for evaluating olfaction against other modalities in multimodal systems, and we believe this to be an important area for future research.

7. ACKNOWLEDGEMENTS

We wish to thank Sidney Fels for his creative ideas and support. Thanks also to Joanna McGrenere for sharing her knowledge of statistical analysis, and to Karon MacLean for recommending we keep our 'fresh' title.

8. REFERENCES

[1] E. Arroyo, T. Selker, and A. Stouffs. Interruptions as multimodal outputs: Which are the less disruptive? In *Proceedings of 4th IEEE International Conference on Multimodal Interfaces (ICMI'02)*, 2002.

[2] E. Cutrell, M.Czerwinski, and E. Horvitz. Notification, disruption and momory: Effects of messaging interruption on memory and performance. In *Proceedings of Human-Computer Interaction–Interact*, pages 263–269, 2001.

[3] M. Czerwinski, E. Cutrell, and E. Horvitz. Instant messaging and interruption: Influence of task type on performance. In *Proceeedings of OZCHI*, 2000.

[4] S. Fels, J. Gauthier, and P. Smith. Responses in light, sound and scent - a therapeutic interactive yoga system. In *Proceedings of the First International Workshop on Entertainment Computing (IWEC'02)*, pages 355–362, May 2002.

[5] S. Greenberg and C. Fitchett. Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of UIST*, 2001.

[6] E. Horvitz and J. Apacible. Learning and reasoning about interruption. In *Proceedings of the Fifth International Conference on Multimodal Interfaces (ICMI'03)*, November 2003.

[7] J. Kaye. Symbolic olfactory display. Master's thesis, Massachusetts Institute of Technology, May 2001.

[8] J. Kaye. Making scents: Aromatic output for hci. In *interactions*, volume 11(1). 2004.

[9] D. S. McCrickard, C. M. Chewar, J. P. Somervell, and A. Ndiwalana. A model for notification systems evaluation - assessing user goals for multitasking activity. In *ACM Transactions on Computer-Human Interaction (TOCHI)*, volume 10, pages 312–338.

[10] B. O'connail and D. Frohlich. Timespace in the workplace: Dealing with interruptions. In *Proceedings of the ACM CHI 95 Conference on Human Factors in Computing Systems*, volume 2, pages 262–263.

[11] M. Weiser and J. S. Brown. The coming of age of a calm technology. www.ubiq.com/hypertext/weiser/acmfuture2endnote.htm.