

The Impact of Unwanted Multimodal Notifications

David Warnock
warnock@dcs.gla.ac.uk

Marilyn McGee-Lennon
mcgeemr@dcs.gla.ac.uk

Stephen Brewster
stephen@dcs.gla.ac.uk

Glasgow Interactive Systems Group, School of Computing Science
University of Glasgow, Glasgow, G12 8QQ, UK

ABSTRACT

Multimodal interaction can be used to make home care technology more effective and appropriate, particularly for people with sensory impairments. Previous work has revealed how disruptive notifications in different modalities are to a home-based task, but has not investigated how disruptive unwanted notifications might be. An experiment was conducted which evaluated the disruptive effects of unwanted notifications when delivered in textual, pictographic, abstract visual, speech, earcon, auditory icon, tactile and olfactory modalities. It was found that for all the modalities tested, both wanted and unwanted notifications produced similar reductions in error rate and task success, independent of modality. The results demonstrate the need to control and limit the number of unwanted notifications delivered in the home and contribute to a large body of work advocating the inclusion of multiple interaction modalities.

Categories and Subject Descriptors

H.5.2 [Information Systems Applications]: User Interfaces—*interaction styles*

Keywords

Multimodal Interaction, Home Care, Reminders

1. INTRODUCTION

A home reminder system is a form of assistive technology designed to help people manage their lifestyle and environment by providing helpful notifications. Wang & Turner [29] argued that home care technology should be adaptable, personalised, customisable and dependable in order to be both effective and appropriate. Many researchers have argued that user acceptability is also a vital factor when developing technology for the home environment [18, 19, 24, 27, 30]. Rejection could lead to the user ignoring or disabling the system [18], limiting its ability to help its users achieve a higher quality of life.

Multimodal interaction can help home care technology to achieve both Wang & Turner's key properties and acceptability requirements by (i) providing interaction opportunities in user-preferred modalities, (ii) providing alternative modes of communications for various social and environmental contexts, (iii) allowing for the consideration of sensory impairment and (iv) providing flexibility in the face of changing requirements. A number of researchers have advocated the inclusion of multiple modalities to make such technology more effective and appropriate [1, 11, 16, 18, 19, 20, 31].

While both hardware and software exists for the inclusion of multiple interaction modalities, a lack of general guidelines and models for their use in commercialised systems has prevented their wide-spread adoption. Existing guidelines for multimodal interaction design need to be refined and validated, which will help give developers the knowledge and confidence to include a wider range of modalities in their technology. In the case of home care technology this could lead to significant improvements in the effectiveness and user acceptability of reminder systems.

The continued development of guidelines and models for multimodal interaction requires wide-scale evaluations of the individual properties of different modalities, both individually and compared with each other. In the case of home reminder systems, this involves examining the underlying properties of different multimodal notifications to determine which ones are more and less effective. Presented here is a study into the potential of unwanted notifications in various modalities to disrupt a home-based primary task. The results of this study will contribute towards the continued development of guidelines and models for multimodal notification design.

2. RELATED WORK

The primary users of home reminder systems are elderly persons and people with sensory or cognitive impairments. Providing multiple ways to interact can improve accessibility, which has prompted a number of researchers to call for the inclusion of multiple modalities in home care technology [11, 19, 25]. In addition, researchers have argued that dynamically switching between modalities when appropriate will also lead to increased user acceptability [1, 18, 30, 31]; for example, switching to a more private modality (e.g. tactile) when delivering sensitive information in a social context might be desirable.

It is well documented that interruptions can cause stress [10, 17], annoyance and anxiety [4], a higher chance of making a mistake [15] and increased speed [7, 9, 17]. Notification schedulers can be used to manage *when* and *how* notifications are delivered to minimise such negative effects. Deciding *when* to deliver notifications has received a great deal of attention in HCI research, mostly focussed around office and high-pressure environments. This usually involves a cost-benefit calculation based on the importance of the notification and how ‘interruptible’ the user is. Horovitz’s attentive user interface paradigm [13] uses attentional focus to measure interruptibility. Monk, Boehm-Davis & Trafton [22] suggested that it was most important to consider the stage of the task; that interruption at the start or end was preferable to interruption mid-task. Iqbal & Bailey’s OASIS [14] framework scheduled around *breakpoints*, naturally occurring mental pauses between units of work, such as when the user has finished writing an e-mail.

When developing notifications for the home context, user acceptability is a key factor. Unhappy or irritated users might sabotage or ignore the system [18], which could be dangerous if it is intended to deliver safety or care-based reminders. There have been numerous attempts at creating a model of when to deliver notifications in the home. Vastenburger, Keyson & Ridder [27] argued that user acceptance of a notification is dependent on the message urgency and the attentional focus of the user, but did not take into account user activity or social context, which other researchers have argued is very important in the home [18, 19, 24]. Nagel, Hudson & Abowd [24] found that people had different opinions on which activities could be acceptably interrupted. McGee-Lennon, Wolters & Brewster [19] reported that people expressed a strong desire for personalisation in home reminder technology.

When it comes to *how* to deliver a notification, there is a great deal of research into different ways to deliver information, but little on how this relates to notification acceptability. Vastenburger, Keyson & Ridder [27] found that the salience of the delivery method combined with the urgency of the notification had a strong relationship to acceptability in the home. McBryan & Gray [18] suggested a configurable model for changing the modality of a notification based around the user’s current activity. McGee-Lennon, Wolters & Brewster [19] highlighted the importance of personal preference when examining modality, and found that interaction modality preference varied with the type of task being carried out. McGee-Lennon, Wolters & McBryan [20] examined 3 types of audio reminders and found that a person’s preferences did not necessarily match the modalities they performed best with.

For the successful inclusion of multiple modalities in home reminder systems, models need to be created that allow developers to identify which modality is appropriate based on factors such as current user activity, user location, social context, user impairment, message urgency and user preference. This requires a good understanding of the relationships between the individual properties of a wide range of interaction modalities.

Warnock, McGee-Lennon & Brewster [31] carried out a study

into the performance and disruptiveness of visual, audio, tactile and olfactory notifications and found that modality did not affect error or activity rates, but did provide different response accuracies and timings. In their experiment, all of the notifications required a response.

However, it is likely that a home reminder system will sometimes deliver notifications that the user will not want or need. Vastenburger, Keyson & Ridder [28] reported that people did not want to be interrupted for low-urgency information. A notification may be unwanted because the user is engaged in an important or demanding task and does not want to be distracted. The notification may not be valid due to a deviation from normal routines: e.g. a reminder to pick up keys before leaving the house may become irritating when the user only wishes to go into the garden. For other occupants of the home all the notifications may be unwanted. It is important to understand the effect that such unwanted notifications could have as an unwanted notification provides no benefit. The inherent differences between modalities may make some modalities harder to ignore or process than others, causing more disruption. This paper reports on an experiment to examine the effects of unwanted notifications in various interaction modalities on performance in a simple card matching game.

3. DESIGN & METHOD

An experiment was carried out to objectively evaluate how disruptive notifications in different modalities were. The experiment was a within-subjects design ($N = 18$) consisting of a primary task and a range of unimodal notifications which instructed the participant to carry out a brief secondary task.

The independent variables were:

1. The modality of the notification (text, pictogram, abstract visual, speech, earcon, auditory icon, tactile and olfactory modalities).
2. The percentage of notifications that were unwanted (100%, 66%, 33%).
3. The type of notification (wanted vs. unwanted).

The dependent variables were:

1. The error rate in the primary task (measured as superfluous views per click).
2. The activity rate in the primary task (measured as clicks per second).
3. The success rate in the primary task (measured as matches per game).
4. The pause time post-notification (measured in seconds).

The experiment was carried out with 18 participants (12 male and 6 female). The participants consisted of 17 people aged 18-30 and 1 person aged 31-45.

3.1 Primary Task

The task chosen was a simple card-matching game called ‘Concentration’ (also known as Memory or Pairs), as used by Warnock, McGee-Lennon & Brewster [31]. In concentration, pairs of cards are presented face-down to the player. The player can then turn over two cards per turn in an attempt to find the pairs and remove them from the game.

Table 1: Modal Specification

Modality	Message			Description
	Heating	Lights	Telephone	
Text	“Heating”	“Lights”	“Phone”	A simple one-word message displayed in a large bold font above the game area on-screen.
Pictogram	IEC-60878 Thermometer	IEC-60878 Light	ISO-7001 Telephone	Taken from two international standards; IEC-60878 and ISO-7001.
Abstract Visual	Yellow Light	Green Light	Blue Light	Projector used to shine a coloured light against the wall.
Voice	Spoken “Heating”	Spoken “Lights”	Spoken “Phone”	Created using the same synthetic voice that was used by McGee-Lennon, Wolters & McBryan [20].
Earcon	Rollercoaster	Static Ride	Water Ride	The earcons varied in rhythm and instrument and were taken from McGookin & Brewster [21].
Auditory Icon	Gas Ignition	Light Switch Click (x2)	Phone Dialing Beeps	Auditory icons at 1 second each taken from an online sound effect archive. ⁴
Tactile	multiSP	textMP	voiceLP	Tactons varied in rhythm and roughness, taken from an experiment by Brewster & Brown [5].
Olfactory	‘Dark Chocolate’	‘Riverside’	‘Raspberry’	Smells were created by Dale Air and selected based on the findings of Brewster, McGookin & Miller [6].

Concentration is a simple leisure activity that might well be carried out at home, it is a well-known game with simple rules and it can quickly build a mental workload. The game was configured in the same way as by Warnock, McGee-Lennon & Brewster; the cards showed simple alphabet caricatures¹ and each game comprised 24 cards with a 60 second time limit.

3.2 Secondary Task

The secondary task in the experiment was to press a single, large red button when instructed to do so by a notification. This button was placed directly in front of the participant within easy reach, as shown in Figure 1. There were two different *notification functions* in the experiment, defined as follows:

The Target Notification

When the participant receives the target notification, they are to press the red button.

Distractor Notifications

When the participant receives a distractor notification, they are to attempt to ignore it entirely and continue focussing on the primary task.

3.3 Notification Design

To evaluate the differences between notifications in different modalities, a wide range of unimodal notifications were designed for the experiment based on those used by Warnock,

¹Speech Teach UK, <http://www.speechteach.co.uk>



Figure 1: The configuration of the experiment.

McGee-Lennon & Brewster [31]. These included common notification techniques such as text and speech along with less common notification modalities such as olfaction and abstract visual display. Eight unimodal notifications in total were developed for the experiment. The exact configuration of each notification is shown in Table 1.

The text and pictogram notifications were delivered directly into the game window to the top of the play area. No additional hardware was required for this. The abstract visual display was created with a short-throw projector positioned to project a coloured light against the wall adjacent to the

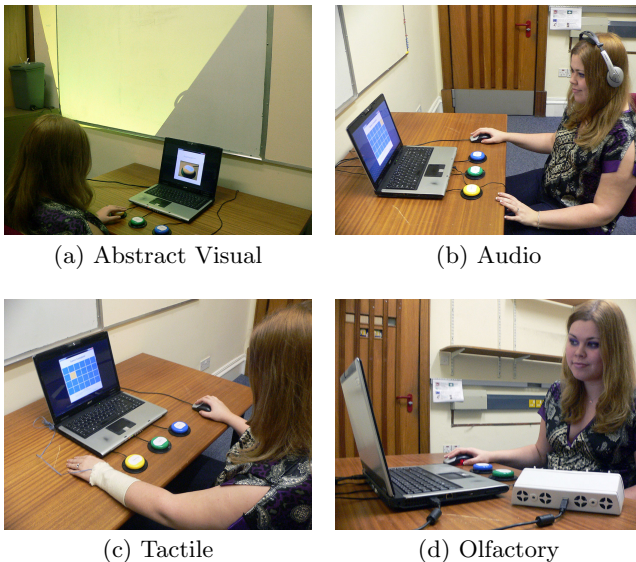


Figure 2: Hardware configurations.

participant. The projector was deliberately aligned so that the projection lay in the peripheral vision of the participant as shown in Figure 2a.

In all audio conditions notifications were delivered through a pair of Sennheiser HD 25-1 II closed-back headphones as shown in Figure 2b. These headphones helped to prevent background noise from causing interference.

Tactile notifications were delivered via an Engineering Acoustics Inc. C2 vibrotactile actuator² powered by a small amplifier. This was secured to the top of the wrist on the participant’s non-dominant hand with a stretchable bandage, as shown in Figure 2c. The device had a very low latency and was able to create precise tactile messages.

Olfactory notifications were delivered using a Dale Air³ Vortex Active smell device, which has the capacity for 4 different smells. Smells are stored on 1-inch disks, which are blown on by a fan to deliver the smell. Delivery times are much longer than other devices, so to ensure the smells were delivered in a reasonable amount of time, the smell device was placed directly in front of the participants as shown in Figure 2d.

3.4 Measurements

Other experiments into disruption have examined errors introduced into the primary task and changes in activity rate post-notification [1, 2, 16, 31]. Latorella [16] observed that activity rate and error rate disruptions have distinctly different properties. In order to measure the effects of a distraction, both error rate and activity rate should be measured.

Warnock, McGee-Lennon & Brewster [31] measured both error rate and activity rate in their card-matching game; activity rate was measured as ‘turns per second’, while error rate was measured as ‘superfluous views per turn’. A super-

fluous view occurs when the participant repeatedly views a card without successfully matching it, suggesting that their mental mapping of cards to locations has been damaged. This type of error more accurately reflects disruption than other error measures (such as failed match attempts) which could be disproportionately influenced by chance or participant playing speed. In this experiment, error rate is measured as ‘superfluous views per click’.

Higher error and activity rates might not necessarily mean a higher failure rate in the primary task. Success rate in the primary task is measured by the number of matches made per game, representing a real-world measure of performance. As discussed in Section 3.1, there are 24 cards meaning this measure has a range of 0 to 12.

The pause time post-notification is measured as the longest event-to-event time in seconds between the start of notification delivery and the end of the notification response period. It is assumed that the longest gap at this point represents the time taken to process notifications when they arrive; for the *target notification* this includes the time taken to receive the notification, process it, stop the primary task, press the button, and resume. For the *distractor notifications* this represents the time taken to receive the notification, process it, and return to the primary task.

Measurements were taken over two ranges. Primary measurements are taken over the course of the whole game, while post-notification measurements are measured from notification delivery till the start of the next notification delivery (or the end of the game). Post-notification measurements are derived from the primary measurements but allow the removal of the unwanted data measured before the first notification, and also allow each notification to be examined individually. As described in Section 3.6, notification delivery times are randomised and split into ‘delivery windows’, meaning that measurements taken post-notification are not taken over a uniform time range; this does not affect the quality of measurement as post-notification measures are expressed independent of time.

3.5 Hypotheses

In order to explore the impact of distractor notifications on the primary task, the following hypotheses were tested.

H1: Notification function will have an effect on the error rate in the primary task post-notification (*measured as superfluous views per click*) and there will be an interaction effect between notification function and modality.

Warnock, McGee-Lennon & Brewster [31] have shown that these modalities produced equivalent changes in error rate when all the notifications were useful, and suggested that it was the act of task-switching that primarily caused the disruption. However, Sanders & Baron [26] found that distractions could display similar properties to interruptions. It is hypothesised that an unwanted notification will lead to a lower error rate in the primary task post-notification as task-switching will not take place. In addition, Warnock, McGee-Lennon & Brewster [31] found differences in response

²Engineering Acoustics, <http://www.eaiinfo.com>

³Dale Air, <http://www.daleair.com>

⁴PacDV Sound Effects, <http://www.pacdv.com/sounds/>

time that suggested some modalities took longer to process than others. This could result in a different levels of error rate change based on the modality of the notification.

H2: Notification function will have an effect on the activity rate in the primary task post-notification (*measured as clicks per second*) and there will be an interaction effect between notification function and modality.

Warnock, McGee-Lennon & Brewster [31] found that modality affected activity rate, which was in line with other research that showed activity rate in a task increases when multitasking [15]. It is hypothesised that the activity rate will not differ with the notification function, as the participants will not know *a priori* if the notifications are distractors or not.

H3: The number of distractor notifications present will have an effect on the success rate in the primary task (*measured as number of matches*) and there will be an interaction effect between the number of distractors present and notification modality.

While it is not expected that notification modality will affect success rate, it is anticipated that distractors will cause less disruption than target notifications, which should mean a higher success rate in games with a higher percentage of distractors. As H1 hypothesises that disruption will vary with notification modality and notification function, it follows that there is likely to also be an interaction effect between the number of distractors present and the modality.

H4: Both notification function and notification modality will have an effect on the pause time post-notification (*measured in seconds*) and there will be an interaction effect between notification function and modality.

Warnock, McGee-Lennon & Brewster [31] found that reaction times for olfactory and tactile notifications were much longer than for other notification types. Although it is hard to differentiate between the mental processing time and actual delivery time, the longer response time for tactile suggested that participants spent longer processing tactile signals. It is anticipated that the different notification modalities will produce different pause times, despite having similar payload delivery times (as discussed in Section 3.3). As target notifications require carrying out a secondary task, it is expected that the pause times for target notifications will be significantly longer, and that there will be an interaction effect between the notification modality and the notification function.

3.6 Procedure

At the beginning of each trial participants were given an information sheet and consent form, then asked to take a short demographic survey to collect gender and age information. Participants were also asked to self-assess their sensory ability on a 21-point Likert scale. Each participant then had the opportunity to practice the primary task (the card matching game described in Section 3.1).

The experiment itself consisted of 8 experimental conditions (one for each modality, as specified in Section 3.3) and a con-

trol condition (where no notifications were delivered). The conditions were delivered sequentially in a random order until all 9 conditions had been completed.

Each experimental condition consisted of a training segment and 4 games. At the start of each condition a screen described the notification modality and the hardware was configured as required (see Sect. 3.3). Each participant was given the opportunity to make minor comfort adjustments (such as volume) before proceeding to the training segment.

Training began by introducing each of the 3 notifications in turn. One of the notifications was then randomly selected to be the 'target notification' that participants were to respond to. Participants pressed a button to reveal this notification and were asked to press the red button (as shown in Figure 1) in order to signal their understanding. Participants could re-deliver the target notification as many times as they liked before continuing. It was stressed to the participants that they should ignore all other notifications except the target notification. All 3 notifications were delivered again, with the participant expected to acknowledge the target and ignore the distractors to complete the training. Participants who failed this test returned to the previous screen and were asked to re-acknowledge the target notification and repeat the test.

The participant then played four games of concentration, each with 3 notifications. Each of the 4 games had a specific configuration, as follows:

- No target notifications, 3 distractor notifications.
- 1 target notification, 2 distractor notifications.
- 2 target notifications, 1 distractor notification.
- 3 notifications randomly determined to prevent participants deducing a pattern in notification delivery.

The 4 game configurations were played in a random order. Notifications in each game were delivered at random points with buffers between them to prevent overlap. A large buffer at the end of the game helped to ensure that quick participants could not finish before all the notifications had been delivered. At the end of each game participants were provided with the opportunity to rest before the next game started.

Once all the conditions had been completed participants were paid and offered the opportunity to ask questions. The experiment required around 60 minutes per participant.

4. RESULTS

H1: Notification function will have an effect on the error rate in the primary task post-notification (*measured as superfluous views per click*) and there will be an interaction effect between notification function and modality.

A within-subjects two-way ANOVA found no main effect of notification modality on error rate in the primary task ($F(7,119)=1.20, p=0.31$), which agrees with earlier studies [31] as expected. Notification function was found to have a significant main effect on the error rate




Figure 3: Graph showing the error rate in the primary task, grouped by modality and notification function.

($F(1,17)=7.36$, $p<0.05$). No interaction effect was found between the notification function and the notification modality ($F(7,119)=0.29$, $p=0.96$).

These results appear to suggest that the distracting notifications actually cause a *larger* change in error rate than target notifications, which can be seen in Figure 3. This result will be discussed in Section 5.

H2: Notification function will have an effect on the activity rate in the primary task post-notification (*measured as clicks per second*) and there will be an interaction effect between notification function and modality.

A within-subjects two-way ANOVA found no main effect of modality on activity rate in the primary task ($F(7,119)=1.16$, $p=0.33$) and no effect of notification purpose ($F(1,17)=0.0$, $p=0.99$). No interaction effects were found between the two factors ($F(7,119)=1.251$, $p=0.28$). This result does not support the hypothesis.

H3: The number of distractor notifications present will have an effect on the success rate in the primary task (*measured as number of matches*) and there will be an interaction effect between the number of distractors present and notification modality.

A within-subjects two-way ANOVA found a main effect of notification function ($F(2,34)=4.84$, $p>0.05$) on the success rate, but found no main effect of modality on success rate in the primary task ($F(7,119)=1.29$, $p=0.26$). No interaction effects were found between the two factors ($F(14,238)=1.40$, $p=0.15$).

As shown in Figure 4, participants generally performed very well. However this graph also shows that the mean number of matches is lower in games with 100% notifications. This finding adds weight to the result of H1, which seems to show that distractor notifications produced higher error rate than target notifications. This will be discussed in Section 5. Post-Hoc pairwise comparisons with Bonferroni corrections




Figure 4: Graph showing the success rate in the primary task, grouped by modality and the number of notifications present.

found no significant differences between the conditions.

H4: Both notification function and notification modality will have an effect on the pause time post-notification (*measured in seconds*) and there will be an interaction effect between notification function and modality.

A within-subjects two-way ANOVA found a main effect of modality on the pause time post-notification ($F(7,119)=14.53$, $p>0.001$) and a main effect of notification function ($F(1,17)=27.63$, $p<0.001$). The interaction test did not satisfy Mauchly's test of sphericity ($X^2(27) = 61.51$, $p < 0.5$), so the results were corrected using the Greenhouse-Geisser correction ($\epsilon = 0.6$). The results found that an interaction effect existed between the notification function and the modality on the pause time post-notification ($F(3.23,54.86)=2.93$, $p<0.05$).

Post-Hoc pairwise comparisons with Sidak corrections found significant differences between some of the conditions; similarly to the findings of Warnock, McGee-Lennon & Brewster [31], olfactory and tactile notifications seem to produce a significantly longer pause than the other modalities tested, as shown in Table 2. The mean difference in pause time between target and distractor notifications is surprising small at 0.29 seconds.

5. DISCUSSION

The primary finding of the experiment was that the error rate in the primary task will increase irrespective of whether the notifications are wanted or not. This disagrees with the suggestion put forward by Warnock, McGee-Lennon & Brewster [31] that the disruptive effect comes primarily from stopping and restarting the task. Experimental conditions without target notifications did not require the users to carry out a secondary task, but still produced a significantly higher error rate ($F(1,)=11.1$, $p<0.001$). In addition, there appears to be no link between error rate in the primary task and response accuracy, showing that even when the participants had difficulty interpreting the notifications, it did not result




Figure 5: Graph showing the pause time post-notification, grouped by modality and notification function.

in a higher error rate. This suggests that the error rate in the primary task must come from another source.

One possibility is that the error rate is not influenced by stopping and starting the task, but as a result of increased mental demand caused by multi-tasking. Participants were unaware of the number of target notifications per game, so may have treated each game equally. Research has shown that a predictable interruption has a lower negative impact [8], but in this case a belief that an interruption may arrive could be the cause of the error rate increase. Bailey, Konstan & Carlis [4] showed that the presence of notifications increased anxiety and the perceived difficulty, which could be the cause of the increased error rate.

Another possibility is that the participants cannot process the notification while carrying out the primary task, so while physical task-switching does not take place, the mental process of stopping and restarting does free resources for the processing of the notification. Research has shown that even the smallest interruption will introduce a negative effect [23], supporting this theory.

Although the source of the primary task error rate increase is unclear, it is apparent that unnecessary notifications cannot

Table 2: Table showing the p -values from *post-hoc* pairwise comparisons between modality and pause time.

Modality	Olfactory	Tactile
Text	$p < 0.05$	$p < 0.001$
Pictogram	$p < 0.01$	$p < 0.001$
Abstract Visual	$p < 0.01$	$p < 0.05$
Speech	$p = 0.12$	$p < 0.01$
Earcon	$p = 0.06$	$p < 0.01$
Auditory Icon	$p < 0.001$	$p < 0.001$
Tactile	$p = 0.97$	-
Olfactory	-	$p = 0.94$

easily be ignored and will produce negative effects. This highlights the importance of a good notification scheduler, able to manage and control notifications of all modalities so that low-priority notifications are not delivered disruptively and that important notifications are delivered in a timely manner in an appropriate modality. This is an active area of research which has already shown how effective a well-designed notification scheduler can be [3, 12, 28]. However, more research is needed to guide the design of multimodal notification systems for the home.

6. CONCLUSION & FUTURE WORK

The study presented here found that unwanted notifications are detrimental error and success rates, but that notification modality does not seem to play a part in this. Modality itself does not affect any of the primary task performance metrics except for longest pause, which suggests that some modalities (i.e. tactile and olfactory) took longer to process than others. That unwanted notifications have the potential to be more disruptive than notifications requiring a response highlights the need for effective notification scheduling systems. This knowledge will help to guide the development of such notification schedulers as well as contribute to the continued refinement of guidelines and models for the inclusion of multiple modalities.

Development of schedulers, models and guidelines depends on a knowledge of which modalities are appropriate when delivering different kinds of information. Vastenburg, Keyson & Ridder [28] have provided the foundation for this work by evaluating the relationship between message urgency and the salience of the delivery method. Further work is needed in this area in order to evaluate the relationship between modality, salience, message urgency and acceptability.

In conclusion, this study has shown that if unwanted notifications are not controlled they are likely to be highly disruptive in the home. In addition, modality was not found to have an effect on disruption independent of the notification value. This implies that home care developers are free to capitalise on differences in delivery speed, information capacity, message obfuscation and conspicuousness of delivery method between modalities to account for factors such as personal preference, sensory impairment, social context, message urgency and user activity. This study provided new knowledge about the role modality plays in delivering notifications, contributing to a large body of work that advocates the inclusion of multiple modalities with the aim of making technology both more effective and more appropriate.

7. ACKNOWLEDGMENTS

This research was funded by the EPSRC (grant number EP/G069387/1). The authors would like to thank the experiment participants and Miss M. Vernigor for modelling the equipment.

8. REFERENCES

- [1] ARROYO, E., AND SELKER, T. Self-adaptive multimodal-interruption interfaces. In *Int. Conf. on Intelligent User Interfaces* (2003), pp. 6–11.
- [2] ARROYO, E., SELKER, T., AND STOUFFS, A. Interruptions as multimodal outputs: which are the

- less disruptive? In *Proceedings. 4th IEEE Int. Conf. on Multimodal Interfaces* (2002), IEEE Comput. Soc, pp. 479–482.
- [3] ARROYO, E., SULLIVAN, S., AND SELKER, T. CarCoach: a polite and effective driving coach. In *Conf. on Human Factors in Computing Systems* (2006).
- [4] BAILEY, B., KONSTAN, J., AND CARLIS, J. The effects of interruptions on task performance, annoyance, and anxiety in the user interface. In *Proc. of INTERACT* (2001), vol. 1, Citeseer, pp. 593–601.
- [5] BREWSTER, S., AND BROWN, L. M. Tactons: structured tactile messages for non-visual information display. In *ACM Int. Conf. Proc. Series; Vol. 53* (Darlinghurst, Australia, 2004), pp. 15–23.
- [6] BREWSTER, S., MCGOOKIN, D., AND MILLER, C. Olfoto: designing a smell-based interaction. In *Conf. on Human Factors in Computing Systems* (2006), ACM, pp. 653–662.
- [7] BURMISTROV, I., AND LEONOVA, A. Do Interrupted Users Work Faster or Slower? The Micro-analysis of Computerized Text Editing Task. In *HCI: Theory and Practice (Part I) Proc. of HCI Int.* (2003), vol. 1, Citeseer, pp. 621–625.
- [8] CARTON, A. M., AND AIELLO, J. R. Control and Anticipation of Social Interruptions: Reduced Stress and Improved Task Performance. *Journal of Applied Social Psychology* 39, 1 (Jan. 2009), 169–185.
- [9] CELLIER, J.-M., AND EYROLLE, H. Interference between switched tasks. *Ergonomics* 35, 1 (Jan. 1992), 25–36.
- [10] COHEN, S. Aftereffects of stress on human performance and social behavior: A review of research and theory. *Psychological Bulletin* 88, 1 (1980), 82–108.
- [11] EMERY, V. K., EDWARDS, P. J., JACKO, J. A., MOLONEY, K. P., BARNARD, L., KONGNAKORN, T., SAINFORT, F., AND SCOTT, I. U. Toward achieving universal usability for older adults through multimodal feedback. *Proc. 2003 conf. on Universal usability - CUU '03* (2003), 46.
- [12] HORVITZ, E., APACIBLE, J., AND SUBRAMANI, M. Balancing awareness and interruption: Investigation of notification deferral policies. *User Modeling 2005* (2005), 433 – 437.
- [13] HORVITZ, E., JACOBS, A., AND HOVEL, D. Attention-sensitive alerting. *Conf. on Uncertainty and Artificial Intelligence* (1999).
- [14] IQBAL, S. T., AND BAILEY, B. P. Oasis. *ACM Trans. on Computer-Human Interaction* 17, 4 (Dec. 2010), 1–28.
- [15] KAPITSA, M., AND BLINNIKOVA, I. Task performance under influence of interruptions. In *Operator functional state: the assessment and prediction of human performance degradation in complex tasks* (Amsterdam, 2003), no. 02, Ios Pr Inc, pp. 323–329.
- [16] LATORELLA, K. A. Effects of Modality on Interrupted Flight Deck Performance: Implications for Data Link. Tech. rep., NASA Langley, 1998.
- [17] MARK, G., GUDITH, D., AND KLOCKE, U. The cost of interrupted work: more speed and stress. In *Proc. 26th SIGCHI conf. on Human factors in computing systems* (Florence, Italy, 2008), CHI '08, ACM, pp. 107–110.
- [18] MCBRYAN, T., MCGEE-LENNON, M. R., AND GRAY, P. An integrated approach to supporting interaction evolution in home care systems. *PETRA 282* (2008).
- [19] MCGEE-LENNON, M., WOLTERS, M., AND BREWSTER, S. User-Centred Multimodal Reminders for Assistive Living. In *ACM CHI* (Vancouver, CA, 2011).
- [20] MCGEE-LENNON, M. R., WOLTERS, M., AND MCBRYAN, T. Audio Reminders in the Home Environment. In *Proc. 13th Int. Conf. on Auditory Display* (Montreal, Canada, 2007), pp. 437–444.
- [21] MCGOOKIN, D., AND BREWSTER, S. Understanding concurrent earcons: Applying auditory scene analysis principles to concurrent earcon recognition. *ACM Trans. on Applied Perception* 1, 2 (2004), 130–155.
- [22] MONK, C., BOEHM-DAVIS, D., AND JG. The attentional costs of interrupting task performance at various stages. *Human Factors and Ergonomics Society Annual Meeting Proc.* (2002), 1824–1828.
- [23] MONK, C., BOEHM-DAVIS, D., AND TRAFTON, J. Very brief interruptions result in resumption cost. In *Proc. 26th Conf. of the Cognitive Science Society* (2004), CogSci 2004, Citeseer, p. 1606.
- [24] NAGEL, K. S., HUDSON, J. M., AND ABOWD, G. D. Predictors of availability in home life context-mediated communication. *Proc. ACM conference on Computer supported cooperative work* 6, 3 (2004), 497.
- [25] NAUMANN, A. B., WECHSUNG, I., AND HURTIENNE, J. Multimodal interaction: A suitable strategy for including older users? *Interacting with Computers* 22, 6 (Nov. 2010), 465–474.
- [26] SANDERS, G. S., AND BARON, R. S. The motivating effects of distraction on task performance. *Journal of Personality and Social Psychology* 32, 6 (1975), 956–963.
- [27] VASTENBURG, M. H., KEYSON, D. V., AND RIDDER, H. Considerate home notification systems: a field study of acceptability of notifications in the home. *Personal and Ubiquitous Computing* 12, 8 (June 2007), 555–566.
- [28] VASTENBURG, M. H., KEYSON, D. V., AND RIDDER, H. Considerate home notification systems: A user study of acceptability of notifications in a living-room laboratory. *Int. Journal of Human-Computer Studies* 67, 9 (Sept. 2009), 814–826.
- [29] WANG, F., AND TURNER, K. Towards personalised home care systems. *PETRA 282* (2008).
- [30] WARNOCK, D. A Subjective Evaluation of Multimodal Notifications. In *Pervasive Health* (Dublin, Ireland, 2011).
- [31] WARNOCK, D., MCGEE-LENNON, M., AND BREWSTER, S. The Role of Modality in Notification Performance. To Appear In *Interact 2011* (Portugal, 2011).