

Control and Personalization: Younger versus Older Users' Experience of Notifications

Isabelle Janzen*

University of British Columbia

Francesco Vitale†

University of British Columbia

Joanna McGrenere‡

University of British Columbia

ABSTRACT

With the increasing ubiquity of mobile technology, users are more connected than ever. Notifications facilitate prompt connections to friends, family and work, but also distract us from what we're doing. We investigated how older and younger users thought about, interacted with, and personalized their notifications. We took a qualitative approach, conducting semi-structured interviews primed through a notification categorization activity. We interviewed 20 participants with equal numbers of younger (19-30 years old) and older (48-74) adults. We extend and refine previous qualitative work and show that while enjoyment plays a minor role in the experience of notifications, urgency, directness and social closeness are far more important factors, though context remains a nuanced issue. We found that older users especially desired a sense of control over their notifications that was difficult to achieve with current technology. Lastly, we provide information about what "categories" of notifications users perceive and expand how that can be used in new personalization systems. These results lead us to advocate a number of fundamental changes to how notifications are personalized.

Keywords: Notifications, Universal Usability, Older Adults, Personalization

Index Terms: Ubiquitous and mobile computing

1 INTRODUCTION

Mobile devices are more ubiquitous now than ever before and deliver a large number of notifications. While notifications can be helpful (e.g. for emergency alerts), they interrupt workflows and hamper productivity [4]. Thus, notification overload is a problem of great research interest. Academic work has tracked the increasing scale of the problem and investigated ways to improve delivery [25, 29].

Our main focus in this work is understanding the experience of receiving notifications. A key motivator was work by Aranda et al. which suggested a model of this experience based on two main dimensions: enjoyment and time urgency [4]. Aspects of that work were used as the basis for the re-design of notifications in the latest Android OS, Oreo [17]. We attempt to further refine and extend their model to support future design.

We consider notifications to be an implementation of interface interruption through a unified notification center that provides different alert modalities (e.g. sound, vibration, silent). These are most common on mobile phones and tablets, but notification centers have also been added to recent desktop operating systems as well e.g., Windows 10. Notifications are varied but we include anything displayed in the notification center of one of these devices (e.g. update install, text message, email, social media awareness).

A key area of interest within the experience of notifications is understanding how that experience varies for a wide range of users.

*e-mail: bjanzen@cs.ubc.ca

†e-mail: fvitale@cs.ubc.ca

‡e-mail: joanna@cs.ubc.ca

By supporting diverse users we may open up new markets and improve the experience for the general population (similar to curb cuts for wheelchairs allowing bikes to travel more easily). We were specifically curious to see if age would impact the notification user experience. Older adults are more susceptible to interruption and have less familiarity with mobile technology [11, 23]. While older adults' experiences with mobile technology has been studied [21], to our understanding, no research has focused specifically on their use of notifications. Thus, one goal of this study is to investigate how the notification needs of younger and older users differ.

We also wanted to understand how these differences could be leveraged to improve notification user experience. One approach is providing personalization settings. There are several competing techniques for providing such settings. These include user controlled, e.g. through a settings panel, adaptive interfaces where settings are adjusted automatically, or machine-human partnerships where settings are suggested to the user [9, 15]. Choosing between these approaches is a complex trade-off between control and convenience.

Regardless of technique chosen, a notification personalization system will require personalizations to be applied to some "type" of notification understood by the user (e.g. *Facebook* notifications). Erroneous personalizations can be particularly problematic for notifications so it is critical that the user clearly understand which notifications are impacted [25]. Previous work has suggested notification filter rules that we conjecture might be confusing to less technical users [25]. Thus, the second goal of our study is to investigate what broader categories of notifications users perceive and how that might impact their personalization needs.

To address our two high-level goals, we conducted semi-structured interviews with 20 participants around their notification usage and personalization preferences. We split the participants into two groups, 10 were younger (19-30 years old) and 10 were older (48-74). We sought to answer the following research questions:

- How do users perceive the impact of the notifications they receive?
- What strategies do users employ to manage their notifications and why?
- What types of notifications do users (consciously or not) perceive?
- What individual differences occur in the above for users within and between age groups?

The contributions of this work are as follows: We refine previous work to show that enjoyment plays a minor role in the experience of notifications, usually dominated by urgency, directness and social closeness. We reveal that older adults in particular desire a sense of control over their notifications that is difficult to achieve with current systems. Lastly, we suggest categorization schemes of notifications informed by what types of notifications users perceive.

2 RELATED WORK

Designing for interruptions is a longstanding area of research in HCI so we focus on work relevant to notifications and older adults.

2.1 Notifications and Interruption

Designing interruptions to minimize cost (e.g. quick time to resume task or low cognitive load) has been well studied. Studies have investigated what moments in task completion are least costly to interrupt [1, 13] and modelled how task recovery is accomplished

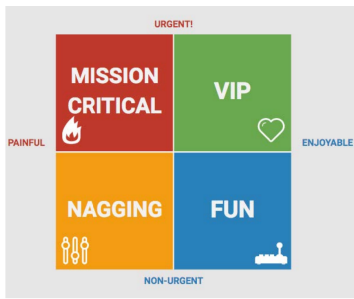


Figure 1: Framework for the experience of Notifications presented by Aranda et al. Several Categories are generated based on varying enjoyment and urgency [4].

[2]. There is also work investigating the impact of different alert design decisions [3] noting trade-offs between alert salience and interruption cost.

Significant work has been done to quantify how often people are interrupted by notifications [19, 29, 32]. Estimates vary, but typically the numbers are considered high and indicative of a significant problem. Complementary studies have investigated qualitative perceptions of what users value in notifications [4, 26]. Users indicate frustration with notification overload and care more about direct messages from socially close contacts than other notifications. There is also work on the etiquette systems users create around notifications [18, 27] noting users associated apps with social groups.

Aranda et al. synthesized this work into the model of notifications displayed in Figure 1 which we attempt to refine [4]. They present two axes, enjoyment and urgency, that determine the experience of receiving a notification. They propose several categories of notifications based on these axes, VIP (for messages from close family and friends) Mission Critical (for highly urgent work) Nagging (for less urgent work) and Fun (non-urgent but interesting).

Much work has tried to improve the user experience of mobile notifications. One such piece of research described PrefMiner, an app that suggests notification filter rules to users, and was a major motivator for our study [25]. Despite their suggestions being conservative, users only accepted 56 percent of them. We conjectured that this might be because it was not obvious what notifications would be removed by the filter rules. Thus, understanding what groups of notifications users perceive that could potentially be personalized around became a major focus of our work. Another direction has been interrupting at opportune times to reduce distraction [13, 28].

2.2 Personalization Schemes

Few users personalize software because of many barriers (e.g. technical knowledge, time investment), but users experience triggers to push through those barriers [5]. To reduce those costs, researchers have developed adaptive interfaces that update settings for users automatically [15]. However, such adaptation comes with complex tradeoffs. If UI elements change position for example, the performance benefit from learning is reduced. Users may also experience frustration if the adaptations are incorrect [15].

Interfaces can incorporate adaptability to different degrees. “Mixed-initiative” personalization, or equivalently “machine-human partnerships” provide greater control to the user. In these systems the computer suggests personalizations to the user and refines those suggestions over time [9, 10]. These systems still have complex design issues. The suggestions themselves may be distracting. Also, the communication must be clear or the user’s response will not be accurate. Our study attempts to understand how users think about notifications so that this communication can be refined.

2.3 Older Adults

Older adults are a growing segment of the population. HCI research involving older adults has used varying age-cutoffs to define “older adult”, e.g., 50, 55, 60, 65, and 70+ [8, 14, 34]. For the purposes of our work, we simply wanted two groups that were relatively younger and older than the other. We aimed for 50+ for our older group, but had two participants just under, as we will explain in the Participants section. Aging is accompanied by declines in perceptual, motor and cognitive abilities. These differences change users’ technology requirements and present tough design challenges [14].

Various research projects have focused on how older adults use and learn new technology [21, 23, 35]. They can face many challenges from simple familiarity with tools and metaphors, to font size readability problems [6]. Many studies have used the new opportunities offered by connected technology to solve life challenges or improve interfaces for older adults [24, 31].

Older adults are more negatively impacted by interruption [11] which poses challenges to the design of certain interfaces and applications [8]. Thus, we conjectured that older adults may be more negatively impacted by notification overload. We also conjectured that older adults may be less able to personalize given the barriers they face learning new technology [23, 35]. However, to our knowledge, no previous study has specifically focused on the notification needs and experiences of older users.

3 METHODOLOGY

Participants completed a semi-structured interview around their notification practices, conducted in a lab. A series of five pilot interviews with members of our lab were done to refine the interview questions and priming activity. After collecting basic demographics, we progressed into the interview which took about an hour. The interview had three stages, with the first two being more structured.

3.1 Procedure

In phase 1, we sought to understand how participants feel about notifications and how they manage them. We aimed to let participants describe these things without pushing them to a specific cognitive frame. We asked open ended, neutrally worded questions about these topics e.g. “How do you feel notifications impact your daily life?”. We allowed participants to talk about whatever the term notifications meant to them, providing our definition if there was confusion. Some participants talked about notifications they got on their tablet, laptop, or phone. We also asked users what strategies they use to manage their notifications.

In phase 2, participants completed a priming activity to get them to reflect on what *types* of notifications they get. Participants were shown a series of 75 cards. Each card had an example notification written on it that covered (1) the app that delivered it, (2) who it’s from and (3) context about the content. Examples include “Email confirming that a package you ordered has been shipped and will be delivered next week” or “Text from your boss telling you to update your timesheet by the end of the week”. After removing notifications each participant thought they would never receive, participants were asked to categorize the notifications in whatever way makes sense to them. Participants then added labels to the categories. We asked follow-up questions around why they created each group, how they might describe it, how they felt about that group, and what notifications within it were particularly important or unimportant.

We used a log of the primary researchers’ notifications over a week to create our initial notification example list. We refined this by asking our pilot participants to share some of their notifications. We found 75 notifications was about the upper limit to keep the activity manageable. We reviewed the examples and balanced them on a couple of coarse subjective factors. These included social closeness (e.g. from a friend, family member or someone you don’t know), and time urgency (e.g. in an hour, tomorrow or next week) which

were important to the experience of notifications in other studies [4]. We made sure several top social media apps had representation, (e.g. Facebook, Twitter, Youtube and Instagram) as well as email, texting and system messages. We tried to make sure that each app had representation for all levels of closeness and urgency, but that was not always possible (e.g. system messages can never be from family). We continued to collect notification examples from participants in the main study, but changes were minor past the first few interviews.

Phase 3 was more open ended. We focused on asking participants how they might like to personalize their devices, now that they had had some time to reflect. Prior to the study we had compiled a rough list of the notification settings available in the default Android operating system, a number of the top-rated apps under a search for “notification managers”, as well as a number of common social media apps. Discussing all settings would have been impossible, so we grouped them into three approaches and asked users how they might feel about them and which they would use. Specifically we asked about: 1. *Removing* unwanted notifications (e.g. filtering, unsubscribing), 2. *highlighting* important notifications (e.g. special sounds, visuals or repeated alarms) or 3. *downplaying* less important notifications (e.g. silent or batched delivery). However, we let participants steer the conversation based on what they wanted, and we brought up nuances and edge cases that had occurred in previous interviews to foster discussion. We also mentioned a hypothetical personalization system that suggested such settings to them, and if they were interested, how much control would they want over it.

3.2 Participants and Recruiting

We recruited two cohorts of ten participants. The first was recruited exclusively from a general call on our university’s paid studies board (four men, six women). Participants simply had to be 18 years old or older, speak English and have used a device that delivers notifications for more than six months. This intentionally skews this cohort to younger adults and was mostly students, a group likely to receive many notifications. The age ranged from 19 to 30 and averaged to 22.9 with Participants coming from a variety of ethnicities. All participants reported significant experience with notifications. They used phones for at least several hours a day and at least periodically checked notifications every hour or two. We refer to this cohort as the “younger” group.

Our second cohort was recruited through the general call on the university board (as per above) and additionally through ads in online classifieds and print ads in the newspaper. In those ads, we specified 50 years of age or older. We did decide to use two participants (48 years old) from the general call even though they were not 50 because (1) they were very near the cutoff, (2) they were much older than the oldest in the younger cohort (30), and (3) their data seemed to align more with the older participants than the younger group. Our goal was having two age cohorts to uncover age effects, but not specifically to map them to particular age ranges. We recognize that this group is a bit younger than what is reported in the older adults literature. Age in this cohort ranged from 48 to 74 and averaged to 58.2 with an equal balance of men and women. Three participants were retired, but still did some contract work. Others had a variety of occupations from computer programmer to fitness instructor. Most regularly used computers as part of their work. Again, participants came from varied ethnicities. They reported similar levels of phone usage to the younger group, but varied more in notification usage. A few reported only checking them once or twice a day (e.g. in the morning and evening). We refer to this cohort as the “older” group.

In this report we refer to younger and older participants by the order they were recruited with an added character to disambiguate the cohort (e.g. Py3 is participant 3, who is younger, and Po8 is participant 8 who is older). Participants in both cohorts were remunerated with fifteen dollars for participating.

3.3 Data Collection

We recorded three data sets. Rough notes were taken by the interviewer, while the groups that participants created in the priming activity were recorded in spreadsheets by the interviewer. The interview was audio recorded and transcribed verbatim.

3.4 Analysis

To analyze our data, we conducted a thematic analysis per the framework presented by Braun and Clark [7]. We chose to allow our themes to be reflective of relatively small groups of participants (even three or four). We expected there to be large variation in participants’ experiences and wanted to find subtle issues that might only impact small groups. We also chose to focus on a descriptive interpretive analysis as we were looking to identify where participants perceived the problems to be. We also attempted to let themes bubble up in an inductive manner because we felt a deductive approach might distract from our goal of understanding user behavior. We note that this research was conducted within a Constructivist philosophical framework. As such we do not claim our results are the objective experience of our participant. They are a rigorously created interpretation necessarily influenced by the analyst’s theoretical experiences and biases.

Analysis progressed in an iterative fashion. Open coding was done by all members of the research team as candidate codes and categories were generated. Coding was collaborative with each researcher being able to see and react to each other’s codes. These were brought to a series of research team meetings to be debated and refined into potential themes. Based on these initial results we iteratively identified additional foci for analysis. Performing additional coding passes and thematic refinement as necessary. For instance, after identifying possible concerns around a lack of control in older participants, we conducted an additional coding pass the younger group to identify any related concerns. During this process the research team continued to meet to ensure the themes continued to reflect the data and remained coherent.

We focused our analysis on reaching *crystallization* [30]. Multiple collaborative coders were used to add complexity to the understanding of the topic through multiple perspectives [33]. We also considered employing member checks, but felt that since the tasks were not related to participants life histories, they might have difficulty remembering the details of their comments months later.

As is common practice in constructivist research, we reflect on how our experiences could influence the analysis. Throughout the study we discussed our experiences with notifications, and problems personalizing them. One researcher strove to continually delete and unsubscribe from unhelpful notifications. Other researchers felt frustrated with how many notifications they got, but couldn’t take the time to experiment with personalizing them more. We acknowledge these experiences are part of what motivates our research goals, and may have informed the issues we focused on in our analysis.

4 FINDINGS AND DISCUSSION

We describe our qualitative data with 3 main themes. Theme 1 (Resigned to Notifications) and Theme 2 (Control and Agency) refer to participants’ responses to questions about how they feel about and manage their notifications. Theme 3 (Types of Notifications) more specifically refers to how participants completed the priming activity and interview.

4.1 Theme 1: Resigned to Notifications

We wanted to get to know the participants and refine previous work around the experience of receiving notifications. In particular, previous work has highlighted how overwhelming notifications can be to users [4]. However, our participants described more nuanced and conflicted feelings. Most had developed strategies to manage notifications that were good enough, if occasionally frustrating. We also

refine the framework for the experience of notifications presented by Aranda et al. [4]; analyzing the nuanced factors involved and suggest that the importance of interesting or fun content in notifications may have been overstated.

4.1.1 Most Participants Have Strategies

The vast majority of our participants had developed strategies they found acceptable to deal with notifications. Most younger participants (eight of ten) had modified notification settings. Most used filtering controls (e.g. filter a group or app) or uninstalled overly talkative apps. A few highlighted important apps or senders by assigning them different alert sounds. Most older participants also used these settings but the proportion was slightly smaller. They also focused more on in-app controls (e.g. unsubscribe buttons). A few put specific focus on email notifications, mentioning categorization rules to organize which previously read notifications they still needed to respond to. No participants in either group reported installing or using a third-party notification manager app.

Despite most participants using notification settings, the majority also used complementary strategies to manage notifications. This was more common in the older cohort who tended to mention these as their primary strategy to manage notifications (8 in 10 versus 6 in 10 for younger). Some participants used *preventative* strategies to prevent “bad” notifications being sent (e.g. selectively installing apps or contact lists). Py10 “I don’t give my phone number to randoms, so, I know that it’s from a friend, family member, my boss, my boyfriend.” Occasionally participants would negotiate *social expectations* of what was sent on each app with senders, (e.g. text me if it’s important) to speed response time. Py3 “If they needed a response...they would text me right away or they would just call me.” About half used *self-discipline* or hiding the phone. Py6 “I try to be more disciplined about it now. Especially when I’m at work.” Several participants also mentioned mental systems to prioritize what they responded to (e.g. respond immediately to emails from the boss, but deal with other notifications later).

4.1.2 Notifications Are a Necessary Annoyance

Participants rarely described notifications as purely helpful, but participants felt they needed to exist. Py3: “They’re disruptive...and I wish that they didn’t exist. That being said, they have to exist in the field of study that I’m in.” We argue many considered notifications more of a necessary annoyance than a major problem. However, these feelings were more negative with some older adults. One older participant (Po8) had had such a negative reaction that he blocked them completely. We unpack this difference more in Theme 2 but continue to examine the broader trend below.

4.1.3 Users Resigned to Notifications

Despite the obvious work to manage notifications, we noticed a common sentiment that the problem was manageable and not worth fixing. This was more prominent in the younger group, but examples came from both. Despite a number of annoyances, many thought that the status quo worked ‘well enough’ not to bother investing more time into improving it.

Py1: “It’s not, I wouldn’t say it’s the nuisance? It’s not too overwhelming at this point, so it’s not a problem, but sometimes it can get annoying, obviously.”

For those who committed to personalizing however, the results could be unsatisfactory or at best bearable. As a case example, take Po17, who put a remarkable amount of effort into personalizing her phone. Po17: “You have to change anything that you don’t want. Which can take hours... I’ve gone from each setting to, you know, in depth to make sure I have everything set up properly.”

Despite this, Po17 was not happy with the end result. Some personalizations didn’t do what she expected them to: “I had set it up so it would alert me, but I didn’t set it up properly so it was

alerting me about everything.....it started telling me about everybody being on their phone [laughs] I don’t need to know that.”

4.1.4 Users: Tell Me What I Need to Respond to

We attempt to refine the framework for the experience of notifications put forward by Aranda et al. [4] We were able to support some of their findings. Specifically that directness, social closeness and urgency are highly valued but we expand by noting that underlying these is whether these notifications have consequences (e.g. missed appointments, ignored friends) for not responding.

In line with Aranda et al., numerous participants expressed that the things they cared most about are direct messages [4]. Py6: “I don’t pay much attention to that [group chats among friends]...unless they specifically message me.” People directly messaging the participant generally expect or desire a response, and might feel rejected if they miss their message.

Social closeness and context also played a large and nuanced role in how important messages from real people generally were. As one might expect, typically closer contacts were more important to respond to because participants cared more and didn’t want them to think they were being ignored. Py9: “These are people, I’m not close friends to, and it’s just very irrelevant.” Po13: “I’d say I have a very close relationship with my immediate family, that’s probably most important, and closely behind that work, then personal.”

However in other cases some younger participants felt friends and family were the people who would understand if you didn’t respond right away. This might be due to our younger participants largely being students at a university with a large international student population where a lag in communication is expected to family in their home country. Py5 (an international student whose parents live in Mexico) showcases this attitude well: Py5: “I would put my parents messages as the least important...I haven’t received shocking news from my parents that I had to address immediately.”

4.1.5 Interested Only if Not Busy

While utilitarian urgency and usefulness of notifications were dominant concerns, participants reacted more negatively to notifications that had uninteresting content (e.g. events they’re not going to). However, how much they valued more interesting but less urgent notifications was heavily contextual. We found that participants generally had little patience for these notifications unless they weren’t busy. Po16: “Maybe for example it’s after 7 or 8 [PM]. And I am interested in somebody’s comments... At that point, very minimally those types of things might drip back into practical.”

4.1.6 Theme 1 Discussion

A key interesting finding in this theme is that current personalizations are being used, but occasional frustration with notification remains. We argue for more research effort to refine these systems and we present approaches to accomplish this in our recommendations.

We also go back to our point that underlying valued notifications is usually whether there are consequences for not responding or attending to the notification promptly. These could be social consequences or utilitarian consequences (e.g. missing an important bill or appointment). This also implied that responding to work messages from one’s boss often took high priority. Po20: “I’d be more impressed or moved if it came from a boss. Only because the need to have a certain impression made.”

Lastly, we find it hard to advocate for interest as a primary determinant of the experience of notifications. We argue that it is essentially the tie breaker: participants value interesting notifications, but not at the cost of interrupting critical tasks. Referring back to Aranda et al.’s [4] framework, we found that the enjoyability axis and “Fun” category seem over emphasized. The VIP, Mission Critical and Nagging categories are consistent with our findings, but are mostly determined by urgency and social closeness.

4.2 Theme 2: Control and Agency

Users have nuanced individual differences that can change software needs and preferences. However, human beings share a variety of near universal basic psychological needs that underlie these preferences [20]. In this theme we mostly refer to the relevant need for a feeling of mastery and competence or self-efficacy [12]. A subset of participants had particularly negative reactions to how notifications interrupted and distracted them. We propose an underlying theme to explain their reactions: that they disrupt users' sense of control and agency over their attention, infringing upon their need for agency and efficacy. We note that these feelings were present in both groups, but more vehemently expressed by our older participants.

4.2.1 Harm from Lack of Control for Some Older Participants

We found a subset of three older participants (Po8, Po13, and Po18) who expressed harshly negative views of how notifications felt disempowering and out of their control. While Po13 had a less extreme reaction than the others, all had made significant life sacrifices in order to regain a sense of control over their attention.

Po8 and Po18 had the most extreme reactions towards notifications. Both participants were older but technically skilled. Po8 was a 48 year old computer programmer. Po18 was a retired 74 year old, who had done mainframe programming. Po8 had tried notifications on his laptop for a while before turning them off entirely because he felt like a "slave" to them. Po18 refused to upgrade to a smartphone to avoid mobile notifications, but still got some notifications on her iPad. Before agreeing to the interview she warned us she only left those on because "I just haven't been bothered to figure out how to remove it - so I put up with it and complain about it." Importantly, these reactions occurred in spite of significant effort to personalize their notifications. Po8 used batched messages for websites that allowed it, and highlighted different senders by ringtones. Po18 had attended several "Facebook for seniors" classes and started unsubscribing from friends sending too many notifications.

Po18 seemed to make a point to avoid notifications because she didn't want to feel dependent upon them, and didn't like how they infringed upon face-to-face time. Po18: "I would say they're an irritant...I refuse to allow myself to become dependent upon them the way a younger person is." She also expressed that she felt like her attention and emotions were being manipulated by notifications, which was harmful to her self worth. Po18: "Sometimes I have total days where I'm not committed, I don't feel like doing anything...and nobody, and the box doesn't ding, then you think nobody likes me [laughs]. It's so irrational and that I see as manipulation."

Po8 however, focused most of his concern on how notifications made him feel like he was no longer in control of his time. Po8: "I look at notifications different than other people. I use social media at your convenience. I don't react to other people's agenda." Po8 describes how the expectations of communication have changed to be more and more immediate. This change may have felt disempowering because it impacted his ability to focus. This may have made him feel unable to complete work effectively and impacted his self-efficacy. This lack of control is referenced in his rejection and removal of notifications "In principle I don't want to react to notifications. We became slaves to our phones." However, this is not to say he doesn't value notifications, just that he doesn't agree to how they are currently delivered. "I'd agree to notifications once a day or twice a day, but once I say yes then I'm at their mercy."

Po13, a 51 year old historian working for the government, mirrored the experiences we noted in Theme 1 and seemed very resigned to the situation. To him, notifications were a fact of life due to his job. However, he felt that he could not stand to be so constantly connected, and forced himself to leave the computer off on Sundays. "They are certainly essential to my daily life... but... I really dislike the thought of being so connected and...I always try to take a break from things on Sunday." This, however, causes some

anxiety about missing important work notifications.

4.2.2 Concerns of Control Are Shared in Both Cohorts

We found that these concerns extended to a small subgroup of younger participants and two more older adults. However, these participants did not make the life sacrifices to address them. The younger group expressed less negativity with this issue, but the underlying sentiment was the same. Py11 made comments that mirrored the more extreme reactions of Po8 and Po18, but laughed about it as a funny annoyance. Py11: "I feel like a slave to notifications at times [laughs]...but, at the same time, I'm grateful for notifications because...notifications help communication feel current."

Py9 also described how she didn't like how the new version of the outlook email client took control away from her and decided to highlight certain emails as important for her. Py9: "They're choosing emails that they think are more important...but I consider all my emails kind of equally important." Similarly some older adults felt annoyed because the alert sound made them pay attention to things they didn't want to. Po4: "Weather apps wouldn't be [so] bad if [they were] silent...What I don't like is the beepy sound."

More commonly, however, a few participants in both groups felt that they didn't want to be on their phones as much as they were, but the demands of work and their social circles forced them to. Some even said that they viewed the instant communication demanded of them as unhealthy. Py10: "I think my usage has spiked because of my work...and it's not healthy."

4.2.3 Desire for Partial Control

This desire for a sense of control has significant impacts on how participants wanted to personalize their devices. All but one participant in both groups expressed that they would like some control in validating or implementing personalizations even if a system made it more automatic. Py6: "I would say that I would still like to have fifty percent control of overseeing what's been filtered through."

Some participants were strongly against total system control Po15: "Probably not. I wouldn't intuitively respond to someone else organizing me" but reacted more positively when we clarified that they could have control to approve suggestions Po15: "Well, suggestions is okay, not so bad."

Interestingly, while we were assessing desired personalizations, we noticed older participants expressed less interest for many basic changes (e.g. filtering). We posit that our suggestions were possibly threatening their sense of control of their attention and knowledge of what they receive. Many older participants wanted to know everything they get so they got to make the decision to respond by themselves. Po13: "To my mind it's either spam or it isn't. If it isn't spam... then I want to decide."

A tension with this is that participants were not always conscious of how many notifications they get, so they would not always have the information they needed to make informed decisions. We noticed participants occasionally being surprised after reading through our example notifications and realizing they receive more than they thought. Po19: "boy I get a lot of notifications [laughs] I can't ignore those." Participants also occasionally re-organized their categorizations during the interview. With even just a bit of probing participants would find edge case notifications they hadn't thought about that might require a different response.

4.2.4 Theme 2 Discussion

We might suggest a few reasons for why younger participants' perceptions of control were less negative. One component may be physiological. Younger participants have been shown to be faster at recovering from interruption [11]. Experience may also play a part: they will have had more time to develop coping strategies. Younger participants may also not have experienced a time when communication was less instant, which would make it difficult to

Table 1: We assigned participants' groupings to the categories on the left through thematic analysis. Counts and example labels are provided for younger (N=10) and older (N=9) participants. One older participant was excluded (Po8) because he had already turned off notifications.

Our Coded Category	Younger Labels	# of Younger	Older Labels	# of Older
Day to Day Reminders	“Util” “Reminders”	8	“Practical” “Calendar”	5
Family and (close) Friends Messaging	“Friends/Family Connections”	5	“Family” “Text”	4
Work	“Business”	7	“Employment”	6
General Messaging	“Facebook Messenger” “Texts”	6	“Personal” “Text”	5
Social Media Activity	“Instagram” “Facebook Posts”	6	“Annoying” “Facebook”	4
System	“Phone System”	6	“System”	1
Email	“Email”	7	“Email”	3

imagine things being different. We also might conjecture that some participants who were not conscious of how many notifications they receive might feel more annoyed if they reflected on it more deeply, though this would need to be investigated in future work.

A key observation to explain the negative feelings in the older group is that notification settings are also often very limited. Generally only allowing users to express a rough idea of “what” (e.g. disabling apps or groups) they want to be notified about, “how” they can be notified about it (e.g. sound or vibrate) and occasionally some contact specific requests (e.g. different ring tones). However, users desiring to express more subtle desires are dis-empowered. Very rarely can users express “when” they want to be notified about something (e.g. a batch every few hours or only in the evening), “why” something might be important enough to overrule these settings or “where” they would want to receive notifications (e.g. only while at home). Lacking the “when” control in particular is problematic as checking email less frequently may reduce stress [22]. To some extent users can adapt their usage of the device to implement these desires (e.g. leaving the phone muted and checking infrequently) but explicit controls should be available.

We argue that any project or technology attempting to improve notification systems must pay special attention to providing users a sense of agency and control over their attention. We note two important design areas that will be impacted: which personalizations are provided and how they are applied. Our recommendations section will provide some preliminary ideas on how these systems can be improved to provide a better sense of agency.

4.3 Theme 3: Complex Categories of Notifications

Previous work has used simple filtering rules (e.g. filter if notification contains “candy” and “crush” in the content) [25] to apply personalizations to notifications, however, they had a low adoption rate [25]. We speculate this was because they might be confusing to non-technical users. It is not immediately clear how many and which notifications the rules will filter. Thus, in this theme we present an overview of the common categories of notifications perceived by users, and how the nuanced contexts surrounding these categories impact personalization behavior.

In the priming activity, participants categorized and discussed a set of 75 example notifications. Through thematic analysis of that discussion we produced our own categorization of notifications seen in Table 1. These are broad categories that tended to underlie the nuanced and contextual groups participants created. For example, a group labelled “text” could be in family messaging and/or general messaging depending on who the participant texted with.

We noticed that a greater proportion of younger participants focused on naming lots of individual apps (e.g. Instagram, Youtube, Reddit) while older participants focused more on describing what aspect of their life the notification related to (e.g. work, family, personal). In particular, the more app based categories (e.g. system and email) were much less common in the older group. Older participants also had more variation in their grouping schemes, and there were a few completely unique categorizations (e.g. participant

16 who simply categorized them by usefulness as 1. Practical 2. Annoying or 3. Don’t care). However, even in the younger group there was large variation in their groupings, and the majority were not just a listing by app name.

The construction of these categories has several implications for notification personalization. For the several system-generated groups like reminders and social media activity, participants’ reactions to them were mostly based on urgency. If a notification from these groups did not provide immediate utilitarian value (e.g. reminding of an important forgotten meeting, debugging a wifi issue, a disaster alert), participants generally did not like or want them. A number of participants also expressed that batching or silent delivery for the less urgent of these would be nice. Py10: “Stuff that I’ll see on my news feed I don’t need a notification for.”

Another important observation was the distinction between several types of messaging including family, personal and work messaging. Similar to the VIP channel presented in Aranda et al.’s model [4], messages from family and close friends are often considered as separate and more important than other messages. However, this did not necessarily mean that responding to them was urgent, just that they needed to see them to decide whether or not to respond. Some friends were considered closer than others and edged into more general messaging. Work messages from bosses also tended to be highly important.

Finally, we noticed that many, especially younger, participants also included a catch-all group for the email they received. However, email is often an enabler for other groups in our categorization. Many events send both a mobile notification and email, which generates a second notification (e.g. a comment response). Despite this, we create a catch-all email group because email often seemed to be prioritized in a very different way. Participants expressed that they used email for more long form communication where immediate responses were not needed. Despite the lack of time urgency, they wanted them quickly because they were being used to discuss more formal topics (e.g. work or school). Py12: “If I’m waiting for an email back from a job...Those all tend to be emails.”

4.3.1 Theme 3 Discussion

We stress that these results are some broad suggestions for common groups and desires in personalized notification delivery. While there are general trends for how participants want to respond to each group, there is almost always significant variation. A social media activity notification could be about some emergency important enough to override users’ original preferences not to be interrupted. Work might have an emergency worthy of a notification over the weekend. Detecting and identifying such nuanced circumstances remains an open and difficult problem.

An interesting age related difference is the decreased presence of email groups from older participants. This may suggest older participants viewed email as a more universal form of communication, whereas younger participants segmented it from day to day messaging; however this would need to be confirmed in future work.

These results prompt us to review how users negotiate what they

are notified about. There should be transparency in what notifications are affected by personalizations. However, a user often agrees to be notified implicitly as part of installing an app. During setup, apps will request permissions from the user (e.g. maps, contacts etc.), and notifications may be included, or simply on by default. Users may feel pressured to accept these permissions, as the app may not work otherwise. However, these systems rarely provide context for why the app needs those permissions or what they will do with them. With respect to notifications, we are missing details such as what will they notify you about, how frequently, and in what way? In addition, each app having its own settings does not line up with how, especially older participants, view notifications. This necessitates duplicate setup for say, similar chat or social media applications.

Lastly, we note that users' perceptions about what notifications they get and which they want evolves over time. Asking yes/no at installation is not an informed decision. Thus, the system asks for the users' forgiveness of unwanted notifications rather than seeking permission. Not to mention that some bad actors explicitly attempt to make it hard to remove unwanted notifications (e.g. tiny unsubscribe links in email ads). As we noted in Theme 2, many participants also shy away from filtering improper notifications. Missing important messages was seen as a significant risk. While the cost for fixing improper notifications falls disproportionately on the shoulders of users, there is still a cost to developers. The annoyance caused by these notifications is often cited as a reason to uninstall apps [4].

5 RECOMMENDATIONS

In contrast to previous work, most of our participants expressed they were not overwhelmed by notifications but had resigned themselves to a "good enough" solution. However, a classic goal of HCI research is taking frustrating but somewhat usable systems and refining them to be more effective. Thus we present two design directions to improve notifications.

5.1 Standardize Personalization above the App Level

There are numerous apps that send notifications installed on every phone, often with overlapping goals and similar notifications sets (e.g. chat messages). Each app requires effort to learn how to personalize, figuring out where the settings panel is and how to use it. We argue that designers need to rethink how notification personalizations are applied and what role the app plays in personalization. Specifically, that personalization needs to be standardized across apps and applied at a higher level.

Many users engage in a lot of personalization, but still experience frustration with notifications. Even if most apps provide proper settings, one or two missed talkative apps is enough to be frustrating. Herein lies the conflict: most apps need to be properly tailored to the user, but that requires a lot of duplicate setup.

We also note that the problem is unlikely to be solved by notification guidelines. Each app developer has differing levels of skill, expertise and goals. It is common for a few bad actors (or naive developers) to try and drive traffic to their site with extra notifications [4]. A first step would be to pull together the notification settings of similar apps (e.g. chat, social media). These more global settings should still allow in-app personalization, but also quick and general settings for common needs. Ideally these categories would be similar to those we describe in Theme 3. However, we acknowledge our categorization may require nuanced and difficult-to-measure contextual preferences. For example, identifying socially close contacts, time urgent notifications, or identifying open periods to deliver less urgent content may involve difficult machine learning.

5.2 Design for Ongoing Machine-Human Partnership

We argue these settings should be maintained with an ongoing machine-human partnership. Participants may not have enough information to make informed decisions about their notifications

during setup. Especially if those personalizations do not do what users expect. This is connected to our previous observation that participants were occasionally not conscious of how many notifications they were receiving. This would make it even more difficult to make an informed decision about which personalizations they need.

We also note that apps are constantly being updated and changed. For example, very recent major changes to the notification system in Android 8 re-designed the notification experience [17]. Grouping notifications from the same app and highlighting direct messages are good first steps to speeding response time. App developers may now offer new settings (e.g. to customize notification shade). While these updates are positive, not all users will know about them, or want to re-personalize.

So how do we maintain these settings? A fully adaptive interface is unlikely to be received well. Many, especially older participants, will reject such a system because it takes control of their notifications away from them. There is a need for an ongoing machine-human partnership. The user must be in the driver seat and have the final say in any personalization, but they should not have to do all the work. The system should highlight key settings and bring them to the attention of the user in unobtrusive ways. Such a system would be able to respond to new settings, as well as refine them as edge cases are uncovered and users get a better idea of what they want.

Some systems have already made attempts to implement such a paradigm [25]. However, more care needs to be taken in how the suggestions are presented to users. Which categories of notifications will be affected by a suggested personalization should be more clearly explained and perhaps be related to our categorization scheme. Further, we argue that attempts need to be made to reduce perceived risk of accepting suggestions, perhaps through offering trial periods for different suggestions. For example "we notice you don't click on this type of notification very much, do you want to try leaving it on silent for a week and we'll check back in later?"

Design iteration is also needed to evaluate what new settings are feasible beyond repackaging currently implemented settings. Users need to be able to specify more subtle wants and needs for their notifications. For example, they might want to specify when and where they are notified and identify higher closeness contacts, but implementation is tricky. Social media analysis may help identify close contacts but privacy remains an issue [16]. Other possibilities may require excessive input or intrusive tracking of location or audio.

6 LIMITATIONS

We focused on providing deep reflections on a smaller group of participants to aide design ideation, rather than a more generalized sample. It would be valuable to evaluate a larger sample of the population to confirm the extent of concerns with agency. In particular, we suspect the older group may have sub-group differences between retirees and not, but we did not have enough data to confirm this. A series of online surveys or more targeted recruiting might help us tackle this issue. This should also include investigating what factors besides age could influence these feelings such as personality traits or cultural background. Studies could also employ more holistic observation to get a broader understanding of notification categorization.

7 CONCLUSION AND FUTURE WORK

This work investigates how younger and older users perceive and think about notifications. Through thematic analysis of semi structured interviews we refined our ideas of user's experiences with notifications. Contrary to previous research, most users seem resigned to the current situation rather than overwhelmed. Participants experiences with notifications also seemed to be far more strongly linked to urgency, directness and social closeness than enjoyment, though context remains a nuanced and difficult issue. We also identified that older users especially desire a feeling of agency over their

attention that is hard to achieve with current notification systems. Lastly, we described common groupings of notifications that users perceive and noted ways this could contribute to personalization systems. These findings lead us to advocate that the current system of personalizations being provided by individual applications needs to be redesigned into a more unified solution. We also argue that an ongoing machine-human partnership will be critical to refining notification personalization. Designers should also look for ways to empower users to specify more nuanced desires for notifications such as when and where they want them to be delivered, as well as identify close social contacts.

REFERENCES

- [1] P. D. Adamczyk and B. P. Bailey. If not now, when?: the effects of interruption at different moments within task execution. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 271–278. ACM, 2004.
- [2] E. M. Altmann and J. G. Trafton. Memory for goals: An activation-based model. *Cognitive science*, 26(1):39–83, 2002.
- [3] A. E. Andrews, R. M. Ratwani, and J. G. Trafton. The effect of alert type to an interruption on primary task resumption. In *Proceedings of the HFES Annual Meeting 2009*, 2009.
- [4] J. Aranda, N. Ali-Hasan, and S. Baig. I’m just trying to survive: An ethnographic look at mobile notifications and attention management. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*, pp. 564–574. ACM, 2016.
- [5] N. Banovic, F. Chevalier, T. Grossman, and G. Fitzmaurice. Triggering triggers and burying barriers to customizing software. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2717–2726. ACM, 2012.
- [6] M. Bernard, C. Liao, and M. Mills. Determining the best online font for older adults. *Usability News*, 3(1):1–4, 2001.
- [7] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- [8] M. Brehmer, J. McGrenere, C. Tang, and C. Jacova. Investigating interruptions in the context of computerised cognitive testing for older adults. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2649–2658. ACM, 2012.
- [9] A. Bunt, C. Conati, and J. McGrenere. Supporting interface customization using a mixed-initiative approach. In *Proceedings of the 12th international conference on Intelligent user interfaces*, pp. 92–101. ACM, 2007.
- [10] A. Bunt, J. McGrenere, and C. Conati. Understanding the utility of rationale in a mixed-initiative system for gui customization. In *International Conference on User Modeling*, pp. 147–156. Springer, 2007.
- [11] W. C. Clapp, M. T. Rubens, J. Sabharwal, and A. Gazzaley. Deficit in switching between functional brain networks underlies the impact of multitasking on working memory in older adults. *Proceedings of the National Academy of Sciences*, 108(17):7212–7217, 2011.
- [12] E. L. Deci, R. M. Ryan, M. Gagné, D. R. Leone, J. Usunov, and B. P. Kornazheva. Need satisfaction, motivation, and well-being in the work organizations of a former eastern bloc country: A cross-cultural study of self-determination. *Personality and social psychology bulletin*, 27(8):930–942, 2001.
- [13] J. E. Fischer, C. Greenhalgh, and S. Benford. Investigating episodes of mobile phone activity as indicators of opportune moments to deliver notifications. In *Proceedings of the 13th international conference on human computer interaction with mobile devices and services*, pp. 181–190. ACM, 2011.
- [14] A. D. Fisk, S. J. Czaja, W. A. Rogers, N. Charness, and J. Sharit. *Designing for older adults: Principles and creative human factors approaches*. CRC press, 2009.
- [15] K. Z. Gajos, M. Czerwinski, D. S. Tan, and D. S. Weld. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pp. 201–208. ACM, 2006.
- [16] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 211–220. ACM, 2009.
- [17] J. Hildenbrad. Android oreo will make you love notifications again. <https://www.androidcentral.com/android-o-and-notifications-everything-you-need-know>, 2017. Accessed: 2017-11-10.
- [18] R. Hoyle, S. Das, A. Kapadia, A. J. Lee, and K. Vaniea. Was my message read?: Privacy and signaling on facebook messenger. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 3838–3842. ACM, 2017.
- [19] S. T. Iqbal and E. Horvitz. Notifications and awareness: a field study of alert usage and preferences. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pp. 27–30. ACM, 2010.
- [20] D. T. Kenrick, V. Griskevicius, S. L. Neuberg, and M. Schaller. Renovating the pyramid of needs: Contemporary extensions built upon ancient foundations. *Perspectives on psychological science*, 5(3):292–314, 2010.
- [21] S. Kurniawan. An exploratory study of how older women use mobile phones. In *International Conference on Ubiquitous Computing*, pp. 105–122. Springer, 2006.
- [22] K. Kushlev and E. W. Dunn. Checking email less frequently reduces stress. *Computers in Human Behavior*, 43:220–228, 2015.
- [23] J. C. Lam and M. K. Lee. Digital inclusiveness—longitudinal study of internet adoption by older adults. *Journal of Management Information Systems*, 22(4):177–206, 2006.
- [24] A. R. McNeill, L. Coventry, J. Pywell, and P. Briggs. Privacy considerations when designing social network systems to support successful ageing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 6425–6437, 2017.
- [25] A. Mehrotra, R. Hendley, and M. Musolesi. Prefminer: mining user’s preferences for intelligent mobile notification management. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 1223–1234. ACM, 2016.
- [26] A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley, and M. Musolesi. My phone and me: understanding people’s receptivity to mobile notifications. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 1021–1032. ACM, 2016.
- [27] M. Nouwens, C. F. Griggio, and W. E. Mackay. “whatsapp is for family; messenger is for friends”: Communication places in app ecosystems. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 727–735. ACM, 2017.
- [28] V. Pejovic and M. Musolesi. Interruptme: Designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 897–908. ACM, 2014.
- [29] M. Pielot, K. Church, and R. De Oliveira. An in-situ study of mobile phone notifications. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*, pp. 233–242. ACM, 2014.
- [30] L. Richardson. New writing practices in qualitative research. *Sociology of sport journal*, 17(1):5–20, 2000.
- [31] M. D. Rodríguez, V. M. Gonzalez, J. Favela, and P. C. Santana. Home-based communication system for older adults and their remote family. *Computers in Human Behavior*, 25(3):609–618, 2009.
- [32] A. Sahami Shirazi, N. Henze, T. Dingler, M. Pielot, D. Weber, and A. Schmidt. Large-scale assessment of mobile notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3055–3064. ACM, 2014.
- [33] S. J. Tracy. Qualitative quality: Eight big-tent criteria for excellent qualitative research. *Qualitative inquiry*, 16(10):837–851, 2010.
- [34] S. Trewin, J. T. Richards, V. L. Hanson, D. Sloan, B. E. John, C. Swart, and J. C. Thomas. Understanding the role of age and fluid intelligence in information search. In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility*, pp. 119–126. ACM, 2012.
- [35] N. Wagner, K. Hassanein, and M. Head. Computer use by older adults: A multi-disciplinary review. *Computers in human behavior*, 26(5):870–882, 2010.