“Put it on the Top, I’ll Read it Later”: Investigating Users’ Desired Display Order for Smartphone Notifications

Tzu-Chieh Lin*
National Chiao Tung University
Hsinchu, Taiwan
jeffreylin98.cs05@nctu.edu.tw

Yu-Shao Su
National Chiao Tung University
Hsinchu, Taiwan
yushao.cs05@g2.nctu.edu.tw

Emily H. Yang
University of Michigan
Ann Arbor, MI, USA
emilyang@umich.edu

Yun Han Chen
National Chiao Tung University
Hsinchu, Taiwan
chelseaxd.pn05@nctu.edu.tw

Hao-Ping Lee
National Chiao Tung University
Hsinchu, Taiwan
dimension4.cs03@nctu.edu.tw

Yung-Ju Chang
National Chiao Tung University
Hsinchu, Taiwan
armuro@cs.nctu.edu.tw

ABSTRACT

Smartphone users do not deal with notifications strictly in the order they are displayed, but sometimes read them from the middle, suggesting a mismatch between current systems’ display order and users’ needs. We therefore used mixed methods to investigate 34 smartphone users’ desired notification display order and related it with users’ self-reported order of attendance. Classifying using these two orders as dimensions, we obtained seven types of notifications, which helped us not only highlight the distinct attributes but understand the implied roles of these seven types of notifications, as well as the implied meaning of display orders. This is especially manifested in our identification of three main mismatches between the two orders. Qualitative findings reveal several meanings that participants attached to particular positions when arranging notifications. We offer design implications for notification systems, including calling for two-dimensional notification layout to support the multi-purpose roles of smartphone notifications we identified.

CCS CONCEPTS

• **Human-centered computing → Smartphones**.

KEYWORDS

Notification Order; Notification Display Order; Notification Management; Experience Sampling Method; Mobile Receptivity; Attentiveness

ACM Reference Format:

Tzu-Chieh Lin, Yu-Shao Su, Emily H. Yang, Yun Han Chen, Hao-Ping Lee, and Yung-Ju Chang. 2021. “Put it on the Top, I’ll Read it Later”: Investigating Users’ Desired Display Order for Smartphone Notifications. In *CHI Conference on Human Factors in Computing Systems (CHI ’21), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3411764.3445384

CHI ’21, May 8–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8096-6/21/05…$15.00
https://doi.org/10.1145/3411764.3445384

1 INTRODUCTION

Smartphone users are receiving an increasing array of notifications on their phones [24]. With the aim of alleviating notification overload, many researchers have investigated factors affecting users’ perceptions that notifications are disruptive (e.g. [24]), as well as their actual phone-notification management practices and preferences [21]. On the other hand, a recent study suggests that while users generally read lists of smartphone notifications in a top-down order, they often start reading them from the middle, and otherwise fail to strictly follow the order in which notifications are displayed [32]. While it is commonly held to be an essential design practice to place frequent and important actions and content in prominent places on a user interface – typically, at the top [12] – it has become unclear whether smartphone users’ current notification-attendance behavior reveals a flaw or limitation in their phones’ current notification-sorting mechanisms. To date, the topic of what kinds of notification users would desire to have displayed at the top of their notification drawers has been under-researched; and studies of discrepancies between their desired notification-display orders and the orders in which they attend to notifications remain rare to nonexistent. Filling this gap in our knowledge will usefully inform future phone systems’ sorting of notifications to match their users’ preferences and needs. Specifically, we seek the answers to the following four research questions:

- RQ1 How do smartphone users’ notification-attendance orders differ from their desired notification-display orders?
- RQ2 How, if at all, do the attributes of the notifications users desire to appear at the top (bottom) differ from the attributes of the notifications they attend to quickly (slowly)?
- RQ3 How, if at all, do users’ desire display order and attendance order for notifications differ across their personal activity contexts?

To answer these research questions, we conducted a mixed-methods study with 34 smartphone users, combining experience sampling method (ESM) questionnaires and semi-structured interviews. Below we first discuss related works and the contributions of the current paper. Then we introduce our methods and the highlights of the results.

2 RELATED WORK AND BACKGROUND

2.1 Mobile Notification and Interruptions

Numerous studies have highlighted the negative impacts of notification overload and/or potential ways of mitigating it [1, 30]. For example, ill-timed notifications can be distracting or interrupting [27], and Leiva et al. [15] showed that interruptions can delay task completion by up to four times. Kushlev et al. [13] reported that external interruptions from notifications could increase the signs and symptoms of inattention and hyperactivity. Other negative effects of being overwhelmed by notifications include decreased productivity and slower and more error-prone performance[11, 25, 31].

As well as to reduce such external interruptions to their own tasks, users often try to reduce the disruption that notification alerts may cause to other people in their surroundings, by silencing their phones [5]. However, while Pielot et al. [28] found that disabling notifications could reduce interference with people's work, doing so can also provoke anxiety. Chang et al.[4] also found that alert variation is highly valued, because it allows people to speculate about notification sources. However, Pielot [26] has suggested that the focus should be on managing expectations, rather than changing the perceptual features of the notifications themselves. After all, smartphone users people are already attentive to their phone, on average once every five minutes within their 12 hours of attentiveness [8]. Yet, both their attentiveness and their responsiveness to specific notifications can be influenced by various notification characteristics, as more fully described below.

2.2 Mobile Notification Management

Researchers [16, 20, 21, 29] found that people valued notifications from messaging apps more highly than other notifications. The identities of the individual or corporate senders [2, 14] and users' subjective views of their importance, urgency and usefulness [21] are also found to influence users' notification-related perceptions and actions. The content of notifications also matters. Fischer et al. [10] suggested that the receptivity to notifications was influenced by how interesting, entertaining, relevant, and actionable each message was perceived to be. Visuri et al. [33] reported that prediction of which notifications were unwanted could be improved if individual notification-content preferences were taken into account. Mehrotra et al. [20] found that people tended to accept notifications containing important or useful content, despite the disruption they caused. It was suggested that the perceived importance and urgency of notifications depended on both their content and users' context. By stopping the 45.8% of notifications that were not useful or important to their participants, Mehrotra et al. [19] successfully minimized the remaining notifications' perceived disruptiveness. Do et al. even [9] suggested that devices' automatic recognition of aspects of their users' physical and social contexts could become critically important, by governing how much and what notification content is served to a particular person in any given situation. However, while smartphone users deem a portion of notifications to be unimportant or unwanted, Westermann et al.[35] showed that only a fraction of them consciously manage their notification settings. Turner et al.'s [32] showed that users typically adopted a top-down approach to manage their notification stacks, but did not always do so. They found that the stack size and notification position negatively affect how users manage notifications and prolong phone usage. Norrie et al. [23] found that their participants disliked managing notifications via a notification bar, instead preferring to use desktop pop-ups on a typing task. And Mashhadi et al. [18] showed that, even though people may not immediately attend to a notification, simply receiving it registers as a visual cue that enables them to go back to it in the near future.

In recent years, Android systems starting with Android O [7] have considered ordering notifications not simply in a chronological order. For example, current Android systems "bucket" certain notifications that normally require user interactions in real time (e.g., driving directions, timers) into a dedicated section, "Major Ongoing", with a secured spot at the top of the notifications list. The "People to People"[1] bucket, containing content links to other individuals such as messages, is also given a high priority. The other buckets include "General", for well-timed and informative task reminders, and "By the way": contextually appropriate or informative but non-urgent content such as weather and traffic reports, and promotions [7]. Such a design also echoes Weber et al.'s [34] findings that within the first five positions in a notification stack or drawer, SMS, IM, and phone notifications were generally concentrated in the top three positions. Although such a design has linked position with attending actions, the relationship as well as difference between desired notification display order and attendance order remains unknown.

Beyond the aforementioned related work, the current study make the following five contributions:

- It is the first research on smartphone users' desired notification display order to show a discrepancy between such order and the current display mechanism.
- It proposes a seven-part typology of notifications, classified into two dimensions: desired display order and attendance order, allowing us to observe the attributes of each of the seven types of notification.
- Its qualitative findings reveal several meanings that people tend to attach to particular positions in the notification drawer, and provide explanations for the mismatches between display order and attendance order that we captured.
- It has several short-term and long-term design implications for future notification systems, including calling for a two-dimensional notification layout to support the multi-purpose roles of smartphone notifications.

3 METHODS

3.1 Experience Sampling Method

We used ESM to let participants express 1) how they would prefer specific sets of six sampled notifications to be ordered in their notification drawers, and 2) in what order they would attend to the notifications in each set. To capture these order data, we developed an Android app using the *Android Notification Listener Service API* [2]. Specifically, the research app recorded all notifications that arrived

[1] The official Android Developers publication on Medium People notifications https://medium.com/androiddevelopers/people-notifications-2a2e4fb6ee96
[2] Notification Listener Service API: https://developer.android.com/reference/android/service/notification/NotificationListenerService

on each participant's phone. The research app then triggered an ESM questionnaire when the following three conditions were met: 1) there were at least six notifications received within one hour, 2) those six or more notifications had been generated by at least three different apps, and 3) those three or more apps contains at least one low-frequency app-combination (e.g. Facebook Messenger + Uber Eats, Facebook Messenger + Weather), i.e., the number of times the combination had been sampled and responded was below the 50th percentile. We chose to sample notifications following these rules rather than directly extracting existing notifications from the drawer because we found in our pilot study that the latter resulted in notification sets highly biased toward app-combinations that arrived more frequently (e.g. messaging apps + emails) and at similar times. To ensure high diversity of sampled notifications from different apps, setting these rules was necessary.

After the above three conditions were met, the app chose the latest six notifications from the selected app-combinations, and prompted an ESM questionnaire. Each ESM questionnaire contained five parts (as shown in Figure 1). In the first part, the participants were asked for information about where they had been and what they had been doing 10 minutes earlier. In the next part, the participants was asked to imagine themselves 10 minutes before, and to sort the notifications based on the order they would have attended to them at that time. We asked the participants to treat all questions in the questionnaires as though it was 10 minutes earlier because it allowed us to obtain orders not only from the high-receptivity moments when the questionnaires were issued, but also, at least possibly, some relatively low-receptivity ones as well. We chose a 10-minute threshold as not being so long before that the participants were unlikely to remember their contexts. The notifications were presented in a list, and participants sorted them via drag-'n-drop (Figure 1(a)). We decided to fix the number of notifications at six because this is typically the largest number viewable in a notification drawer without scrolling. Notifications that participants felt they did not need to click on could be "swiped" to label them accordingly, as seen in Figure 1(b). Whenever this happened, a pop-up prompted the participant to select a reason s/he thought it unnecessary to attend to the swiped notification at that earlier moment, from among: 1) "No further information needed"; 2) "No further action needed"; 3) "Not interested"; and 4) "Other.", in the last case, with free-text field into which another reason or reasons could be input.

After they did their sorting, the participants were asked four five-point Likert-scaled questions about each of the six notifications, covering its importance, urgency, content attractiveness, and sender attractiveness, as shown in Figure 1(e). Since the measurement of content attractiveness [3] and sender attractiveness did not make sense in all cases (e.g. a notification not related to a sender), a "Not applicable" option was also provided in the question.

In the next part, the participant sorted the same notifications according to their desired display order (Figure 1(c)). Here, again, they could swipe a notification to express that it did not need to be displayed at all, upon which a pop-up would prompt them to select a reason they thought it unnecessary to display at that earlier moment. These reasons were: 1) "Repetitive"; 2) "Known information"; 3) "Not interested"; and 4) "Other", again with a text box. If our research app detected any difference between someone's attendance order and his/her desired display order, it prompted that person to view the two order lists and explain why the two orders were different (Figure 1(d)).

A minimum duration of 2 hours was interposed between any responded-to ESM prompt and the next such prompt; and the maximum number of questionnaires per day was set at five. To minimize inaccurate self-reporting caused by recall bias, an ESM prompt was dismissed after 10 minutes, but once it expired, the app could trigger another ESM as soon as it detected and recorded six notifications that met the criteria described earlier.

3.2 Recruitment and Participants

We recruited participants via several Taiwanese Facebook groups intended for research-subject recruitment. Of the 34 participants that participated in our study, the youngest was 21 and the oldest, 38 (M=25.53, SD=5.02); 17 were students and 17, non-students; and 20 were male, 14 female. All participated for at least four weeks, with one participant extending the duration of the experiment for three days voluntarily. They had diverse self-reported numbers of notifications per hour: with 12 participants saying they received 2-5; another 12 individuals claiming 6-10; and 10 claiming more than 10. They also had diverse self-reported frequencies of clearing the notification drawer: with 16 cleared more frequently than once an hour; the other 18 cleared less frequently than once an hour. All reported receiving notifications from at least four apps on their phones every day.

3.3 Study Procedure

Each participant was invited to an in-person pre-experiment briefing and a post-study interview. In the first meet-up, the research team explained the study procedure and helped them to install the research app. After the four-week experiment, all the participants were invited to participate in interviews, and 33 accepted. In the interviews, they were presented with a list of apps that our research app had frequently logged, and asked about how they would sort notifications from each, based first on their typical attendance order, and then on their desired display order. This process was used primarily to prompt participants to reflect on their attitudes to the notifications from each app, and the reasons for discrepancies between their two orders (if any). Participants received compensation according to the number of ESM questionnaires they had finished, at the rate of NT$8 (at the time of writing, US$0.27) for each questionnaire. If they finished more than 60 questionnaires, they were given a bonus of NT$100 (US$3.41) for every 20 more additional finished questionnaires. An additional NT$400 (US$13.64) was given to them if they also participated in an interview. This study was approved by our university's Research Ethics Committee for Human Subject Protection.

3.4 Categorizing apps

All of our ESM collected notifications came from 473 unique apps. Two coders independently derived app categories for the 473 apps based on Google Play Store listings and relevant prior literature [30, 34, 36]. Their respective sets of categories were then discussed and finalized, which resulted in a total of 18 categories. Then, the

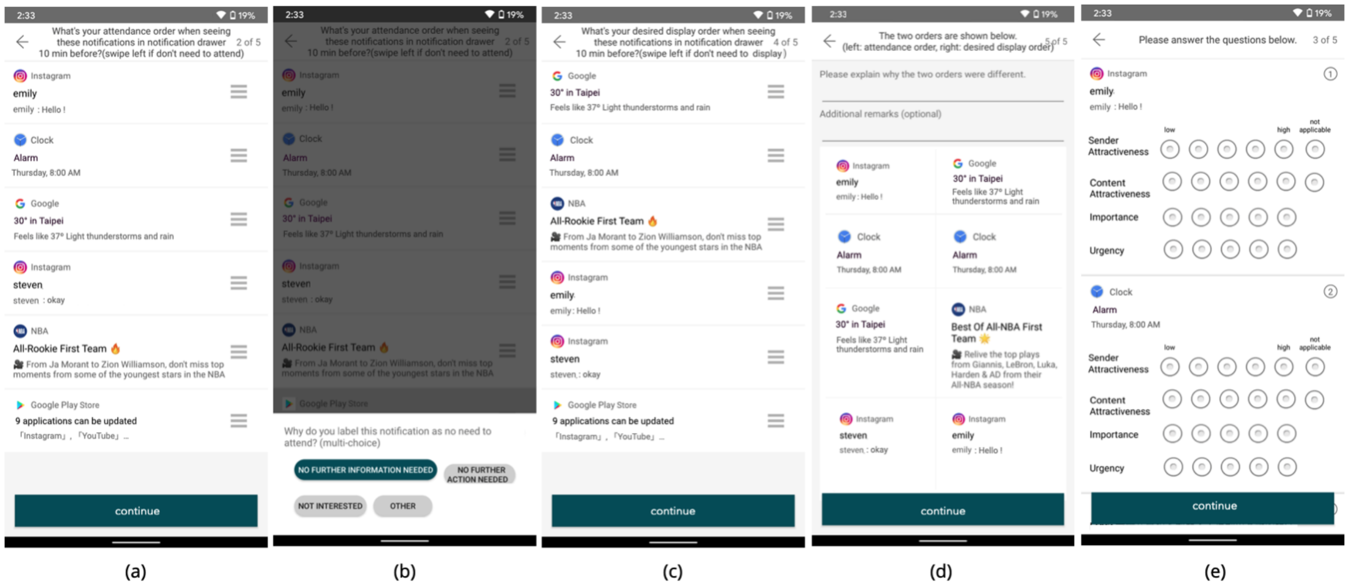

Figure 1: In the ESM questionnaires, (a) Users would sort 6 notifications based on their attendance order. (b) To indicate that they did not want to attend to a notification, they could swipe it away, and a pop-up would ask for reasons. (c) Users would also sort the notifications based on their desired display order. (d) When the two orders were detected to be different, users could provide explanations. (e) Users would provide their perceptions of each sampled notification. For each part, users were asked to imagine themselves seeing the notifications in the drawer 10 minutes before while answering.

same two coders coded all of the apps into one of these 18 categories independently, and achieved an intercoder agreement Cohen's kappa coefficient of 0.839. The final consensus on coded categories of all the 473 apps were made after the discrepancies were discussed.

3.5 Quantitative Data Cleaning and Analysis

We collected a total of 1,952,369 notifications (1,310,881 ongoing notifications, 641,488 non-ongoing notifications) from the 34 participants. The top category was System (36.93%), IM (24.89%), Transportation (20.01%). The number of apps per participant that posted any notifications during the study period ranged from 26 to 119 (M=52.85; SD=19.28). We received a total of 3,335 ESM responses, with a response rate of 42.6%, meaning that a total of 20,010 notifications were sorted. We eliminated any ESM responses that were completed within 30 seconds, on the assumption that the person completing it had not been giving it sufficient effort and attention. We eliminated a total of 219 ESM responses on that basis, and analyzed the remaining 3,116. Moreover, a total of 257 notifications were marked by participants as ones they did not want to be displayed, but which they wanted to attend to anyway. We removed them from analysis due to the difficulty of interpreting what attending to non-displayed notifications would mean in terms of real-world behavior. We yielded a final total of 18,439 user-sorted notifications. Among this final set of sorted notifications, the top five notification categories were System (18.46%), IM (16.13%), Shopping (9.16%), Social (8.66%), Weather (7.16%). There were 170 unique two-app combinations of notification categories, with the top five combinations being IM+System (6.59%), Social+System (3.27%), Shopping+System (3.26%), Weather+System (3.26%), System+System (3.18%).

3.6 Qualitative Data Analysis

We transcribed all interview recordings and subjected them to affinity diagramming [17] with a bottom-up process. Through iterative grouping and labeling, several themes regarding the participants' attendance order and display order emerged. These included: general attendance and display order; mismatch between the orders; the meanings of particular display positions; the influence of context; and attitudes toward current notification systems.

4 QUANTITATIVE RESULTS

4.1 Desired vs. Current Display Orders

To answer our RQ1, we first present the difference between users' desired display order and the current system. We extracted notifications from Phone, IM, Media, Transportation, and Calendar & Reminder apps that were marked as "ongoing notifications" by the system, and deemed them to be *Major Ongoing* ones, resulting in a final pool of 583 Major Ongoing notifications. All IM notifications were labeled as *People to People* ones unless they had already been classed as Major Ongoing. Figure 2 shows the distribution of notifications from each Android-assigned category being placed in each position, which was quite even. 55.3% of the Major Ongoing notifications (Figure 2, left bar) and 56.1% of the People to People notifications (Figure 2, middle bar) were *not* desired to be placed in any of the three highest positions. Only 16% of them were desired to be placed uppermost, while a surprisingly high 29.7% of Major Ongoing notifications were deemed not necessary to show in the

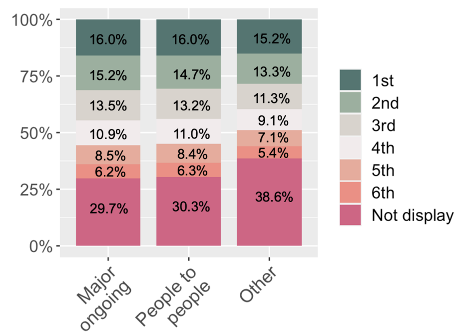

Figure 2: Desired display orders of "Major Ongoing" and "People to People" notifications, as against all notifications not in these two Android-assigned categories.

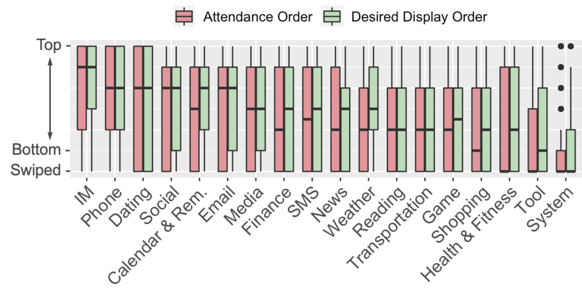

Figure 3: Box plot showing the distribution of desired display order and attendance order for each notification category (bar=interquartile range, median at center line, points represent outliers).

drawer at all. Among that subset of the Major Ongoing notifications, two-fifths were swiped off because participants thought they already knew the information being provided. Among the other 60% of this subset, half were swiped off as repetitive, and the other half as uninteresting. This starkly contradicts the design rationale of the current display mechanism, i.e., that Major Ongoing notifications require users' attention and active awareness because of their high importance [7]. These results also show discrepancies between the current display mechanism and the way participants wanted their notifications to be displayed.

4.2 Attendance Order vs. Display Order

Figure 3 shows the boxplot of the desired display order and attendance order. Top position means that the participants desired to place the notification in the 1st position or want to attend to it first. While the distribution between the two orders was similar in most notification categories, there were some discrepancies between the two orders from some categories. The discrepancy in the distribution of the Weather-related notifications was the largest, and was also larger for Calendar & Reminder. Participants tended to desire to display them at a relatively high position, but would not attend to them soon, or even not at all. News and Tool notifications also had a higher medium of desired display order than the attendance order. Health & Fitness is a particularly interesting case, where the median of its attendance order was zero (i.e., not attended to), while the median of its desired display order is two (i.e., second lowest). This was because participants often swiped them off (56% of the time) because they felt uninterested in them (60%).

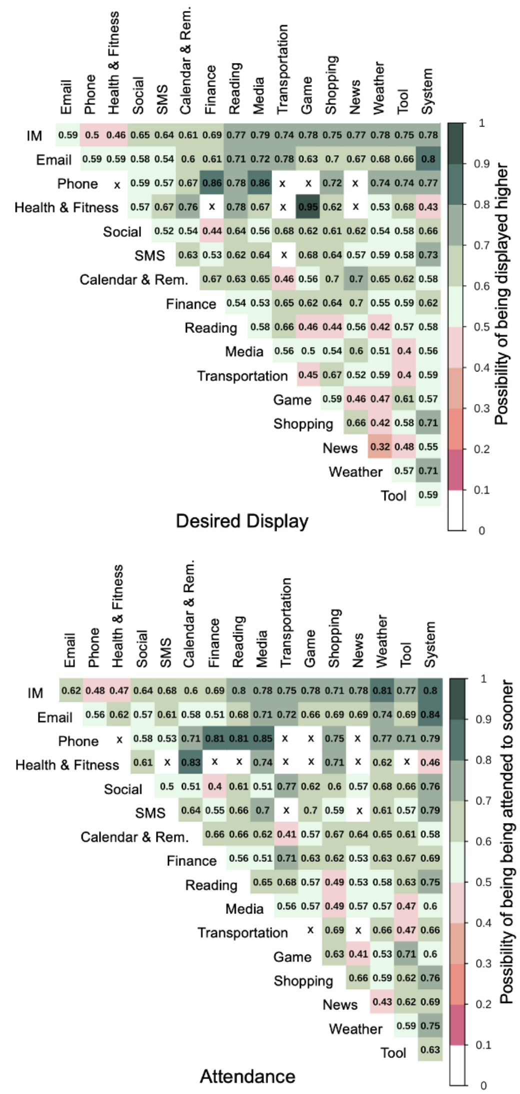

Figure 4: Pairwise duel matrix. The number inside a grid represents the possibility that the category in that row is displayed higher or attended to sooner than the category in that column ("x" represents count lower than 20).

4.3 Pairwise Duels Between App Categories

For notifications that participants did want to display or attend to, we wonder, when two of these notifications, of a different category, were presented concurrently, which one participants would desire to attend to sooner or position higher. We generated all combinations of any two notifications that were sorted together in ESM questionnaires. Figure 4 shows this dueling relationship matrix. We only included pairwise duels of which the count was over 20 in the

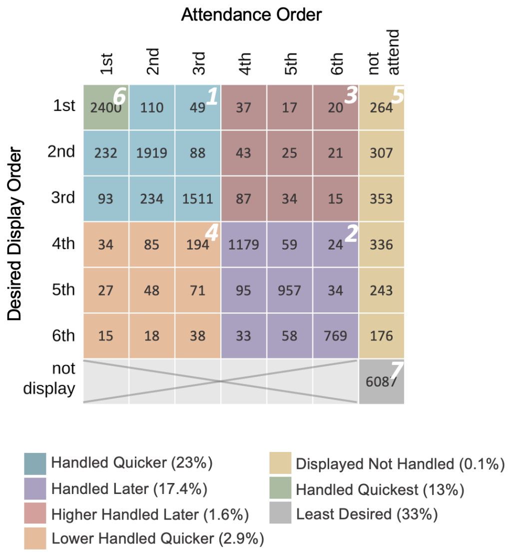

Figure 5: Overview of the seven types of notifications, including three type of mismatches (Block 3, 4, 5). The percentage of each type are also presented.

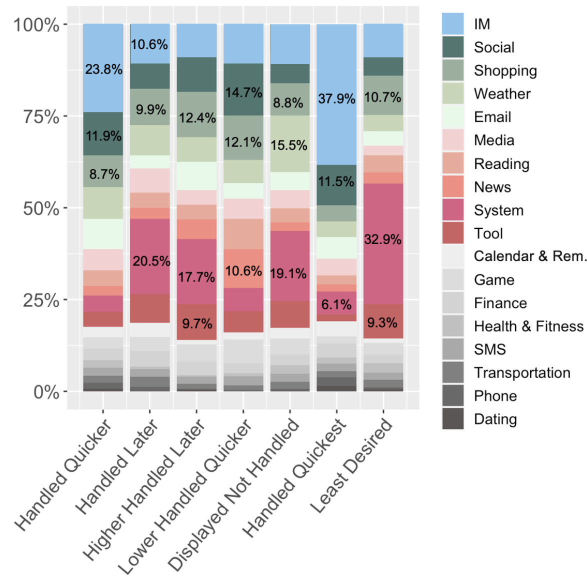

Figure 6: The percentage of each notification category in the seven types of notification. Only the percentage of top three notification categories are labeled.

figure. The figure shows that four categories generally won in pairwise duels in both orders: IM, Email, Phone, and Health & Fitness. Interestingly, Health & Fitness won most duels against the other notification categories. A significant portion of these notifications were swiped out (i.e., not attended to), but we found that when participants needed these notifications in the notification drawer, they more often attended to them faster and wanted to place them higher to fulfill their momentary needs. The top three activity contexts associated with these instances were resting, eating, and moving, resulting in those notifications being more relevant to those specific moments than the other types of notifications.

4.4 Seven-part Typology of Notifications

To answer our RQ2, we classify all notifications using the two orders as dimensions into seven types. Figure 5 shows the number of notifications that participants placed at each of the six positions in our ESM questionnaires, when sorting them for either of the two orders. The ranking from top to bottom represents the desired display order, whereas the left-to-right one is attendance order. Among all the notifications being analyzed (18,439), 67.0% were desired to be displayed (Figure 5, Blocks 1-6). Among these notifications, the majority (9,844 out of 12,352; 79.7%) were placed at the same position in the desired display order and in the attendance order; such congruence will be referred to hereafter as *matches*; and the other 20.3% as *mismatches*. In Figure 5, mismatches are clustered in three locations: Block 3, Block 4, and Block 5. The seven blocks in Figure 5 show the seven types of notifications, of which the quantity, as a proportion of all notifications, are shown. These seven types are: *Handled Quicker (Block 1), Handled Later (Block 2), Higher Handled Later (Block 3), Lower Handled Quicker (Block 4), Displayed Not Handled (Block 5), Handled Quickest (Block 6), and Least Desired (Block 7)*. Figure 6 and Figure 7 show the percentages of each notification category and the average attributes of the notifications for each type respectively. The major notification categories of each type are distinct. A Chi-Square of independence test shows that the distribution of the percentages is strongly related to the type (χ(102, N=18439)=4077, p<.001). We also used a mixed-effect logistic regression model to examine whether the four attributes, *importance*, *urgency*, *sender attractiveness*, and *content attractiveness*, were correlated with each of the seven types. Dummy coding was used for this logistic regression model. Below we discuss the seven types in detail.

4.4.1 Moderately Matched: The Handled Quicker and Handled Later. *Handled Quicker* were notifications participants wanted to see in the top positions and would attend to sooner. They were rated as having relatively higher average importance (2.4), urgency (2.0), sender attractiveness (3.1), and content attractiveness (3.0). The regression result shows positive effects of content attractiveness (Z=6.90, p<.001) and sender attractiveness (Z=2.75, p<.01), suggesting that notifications that were attractive were more likely to be placed in higher positions and attended to quickly. The top app categories in this group are IM (23.8%), Social (11.9%), and Shopping (8.7%).

Handled Later were notifications participants wanted to see in the lower positions and would attend to later. They were rated as having relatively lower average importance (1.7), urgency (1.5), sender attractiveness (2.1), and content attractiveness (2.0). Regression result shows negative effects of importance (Z=-7.50, p<.001) and both sender (Z=-3.16, p<.01) and content attractiveness (Z=-5.05, p<.001), suggesting that notifications of high importance or

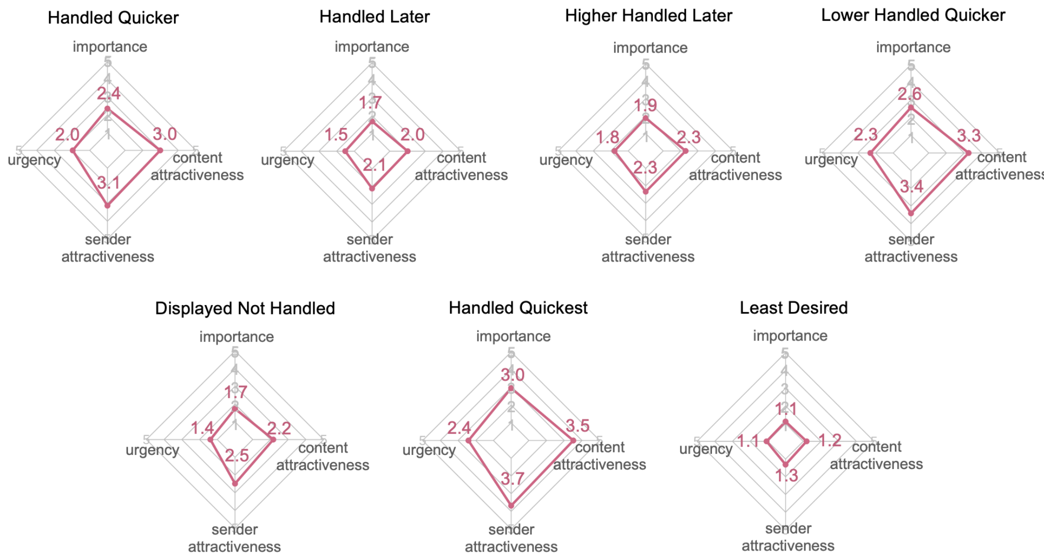

Figure 7: Average ratings of the four attributes: *importance, urgency, sender attractiveness,* and *content attractiveness* for each type of notification.

attraction were less likely to be placed in lower positions and attended to slowly. The top app categories in this group are System (20.5%), IM (10.6%), and Shopping (9.9%).

4.4.2 Three Mismatch Types: Higher Handled Later, Lower Handled Quicker, and Displayed Not Handled. There are three types of mismatches. The first type, referred to as *Higher Handled Later*, were notifications that participants desired to be placed at the top, but would attend to later. They were rated as having relatively lower average importance (1.9), urgency (1.8), sender attractiveness (2.3), and content attractiveness (2.3), but slightly higher than those of *Handled Later*. Regression shows a negative effect of sender attractiveness (Z=-2.51, p<.05) of this type, suggesting that notifications involving senders were less likely to be placed at the top without attending to them soon. The top app categories in this group are System (17.7%) (e.g. low battery warning, software update, Wifi connection), Shopping (12.4%), and Tool (9.7%).

The second type, referred to as *Lower Handled Quicker*, were notifications that participants desired to place at the bottom but wanted to attend to sooner. Compared to the previous mismatch, notifications of this type were rated as having higher average importance (2.6), urgency (2.3), sender attractiveness (3.4), and content attractiveness (3.3). Surprisingly, the four attributes were even higher than those in *Handled Quicker*. This suggests that users do not always want to put attractive and important content in higher positions. Regression result shows negative effects of importance (Z=-3.98, p<.001) and positive effects of both urgency (Z=3.49, p<.001) and content attractiveness (Z=3.13, p<.01). These notifications were generally those considered time-sensitive and needed/interesting content for the participants, which per se were not important to place at the top, but were urgent enough to act fast to obtain the content, such as discounts and promotions with time constraints. The top app categories in this group are Social (14.7%), Shopping (12.1%), and News (10.6%).

The third type, referred to as *Displayed Not Handled*, were notifications that participants desired to display in the notification drawer but did not want to attend to. This is the most common mismatch. From the reasons the participants chose for why they did not want to tap them, 57.0% chose "*Don't need further information*", suggesting that the information supplied in the notification text itself already sufficed. Notifications of this type were rated to have relatively lower average importance (1.7), urgency (1.4), sender attractiveness (2.5), and content attractiveness (2.2). Regression result shows positive effects of importance (Z=4.95, p<.001) and negative effects of urgency (Z=-10.18, p<.001) and content attractiveness (Z=-2.83, p<.01), showing their relatively higher importance compared to Least Desired, but lower urgency and content attractiveness. The top app categories in this group are System (19.1%), Weather (15.5%), and Shopping (8.8%).

4.4.3 The Two Extremes: the Handled Quickest vs. The Least Desired. *Handle Quickest* were notifications participants attended to the soonest and wanted to display at the top. Regression result shows that attributes of all importance (Z=8.19, p<.001), urgency (Z=5.36, p<.001), content attractiveness (Z=4.42, p<.001) and sender/app attractiveness (Z=8.07, p<.001) have positive effects. We particularly separated *Handle Quickest* and *Handle Quicker* because the participants had the strongest desire for these notifications. Notifications of this type were rated as having the highest averages for all attributes: importance (3.0), urgency (2.4), sender attractiveness (3.7), and content attractiveness (3.5). The top app categories in this group are IM (37.9%), Social (11.5%), and System (6.1%).

Least Desired was the largest notification type (nearly one-third of the notifications), which were neither desired to be displayed nor wanted to be attended to. From the reasons participants chose for why they did not want to display them, more than half (55%) of the reasons were: "*I'm not interested*". Indeed, notifications of this type were rated to have the lowest averages for all attributes:

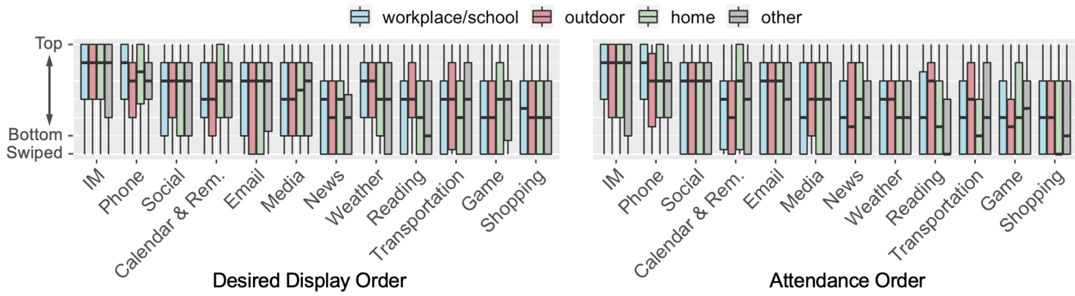

Figure 8: Box plot showing the distribution of desired display order (left) and attendance order (right) for each of the three location context (Home, Workplace, Outdoor) by notification category.

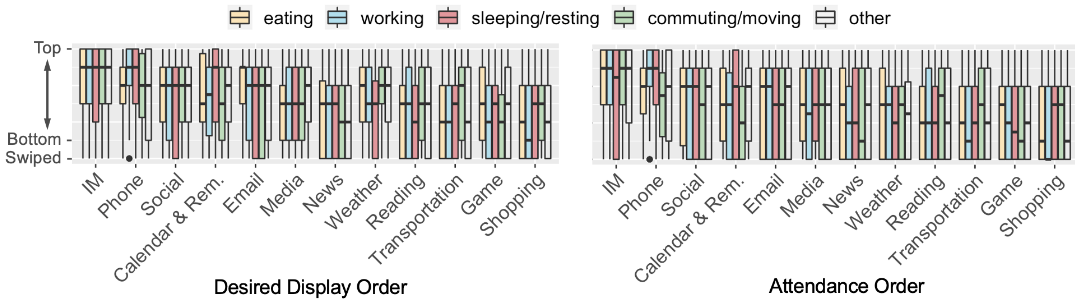

Figure 9: Box plot showing the distribution of desired display order (left) and attendance order (right) for each of the four ongoing activities by notification category.

importance (1.1), urgency (1.1), sender attractiveness (1.3), and content attractiveness (1.2). The top app categories in this group are System (32.9%), Shopping (10.7%), and Tool (9.3%).

To summarize, the seven-part typology allows us to identify different kinds of notifications. In spite of some similarities between certain types, most of them are distinct in terms of either their attributes or representative categories, suggesting their different roles as notifications, especially manifested by the existence of the three mismatches. Furthermore, it also implies the meanings of positions, so that some attractive and fast-attended-to notifications (e.g., *Lower Handled Quicker*) were not placed at the top.

4.5 Desired Display Order across Location and Activity Contexts

Finally, to answer our RQ3, we examined how context influences participants' desired display order. Figure 8 and Figure 9 show participants' desired display order and attendance order across location and activity contexts, which was extracted from participants' responses to the ESM questionnaires. We used ordinal regression to examine whether contexts were correlated with the two orders. Email and Phone notifications were positioned higher when participants were at their workplace/school (M=3.6; SD=2.2) than when they were outdoors (display: M=3.2; SD=2.2). The difference was significant (Z=-2.01, p=.04). Reading and Transportation related notifications were positioned higher and attended to faster outdoors than at home (display: M=2.2;SD =2.2, attendance: M=2.1;SD =2.2) and at the workplace/school (display: M=2.5;SD =2.0, attendance: M=2.3;SD =2.2). The regression result shows that the differences are both significant (outdoor vs. home, display: Z=-3.56, p<.001; attendance: Z=-3.75, p<.001; workplace/school, display: Z=-3.83, p<.001 (display), attendance: Z=-2.67, p<.01). Calendar & Reminder notifications were positioned higher and attended to faster at home (display: M=3.6;SD =2.3, attendance: M=3.4;SD =2.4) than when outdoors (display: M=3.0;SD =2.0, attendance: M=2.5;SD =2.1); the difference was also significant (display: Z=-2.16, p=.02, attendance: Z=-2.39, p=.02).

We also observed differences in the effect of context between the two orders. For example, when eating, participants tended to place Email notifications in higher places (M=3.8;SD =2.0). Our ordinal regression result shows that eating (Z=2.30, p=.02) have effects in predicting the desired display order for Email notifications, but displayed diverse attentiveness to these notifications. These results together suggest that while the influence of context on the desired order and attendance order was similar, for certain categories of notifications the influences were different.

5 QUALITATIVE RESULTS

5.1 How Participants Desired to Display Notifications and Why

Unsurprisingly, given that their impacts on attendance order have frequently been highlighted in notification research, factors including perceived urgency (P24), importance (P34), the sender of the notification (P16), the frequency of using the apps (P32), and personal relevance/interest (P24) were all mentioned by our participants as influencing their desired display order. However, many more factors were mentioned by the participants as influencing their desired display order, and these factors were principal drivers of the observed mismatches. We present these factors below.

5.1.1 Notifications as Reminders. One commonly mentioned reason for desiring a particular notification to occupy a particular position was that the interviewee was treating it as a reminder. Yet, their choices of where such a "reminder" notification ought to be placed varied widely, according to their general notification-handling preferences. Those individuals who perceived themselves as receiving many notifications but did not carefully review them consistently expressed worries that these "reminders" would be "buried" in their notification piles. Thus, they preferred to place such notifications at the top of the drawer: "*I set different reminders at different times of a day. So the thing that I remind myself of is important and needs to be placed on the top*" (P13). Some expressed their hope that such notifications could be "pinned" to the top, so that they could serve as easily noticed visual cues. As P31 explained: "*For those that are important, but that I can't handle right away, I wish there was a way I could pin them. You don't have time now, but when you have some eventually, you'll know it'll be on the top*".

Notifications characterized by the interviewees as "reminders" came from diverse sources; and many were advertisements and/or contained discount information. As P13 commented, "*It's not only about whether it's important or urgent, but about whether I intentionally left it there, so that I wouldn't forget.*" The same participant provided the following further example: "*I have some vouchers I need to spend soon, so I'll pay more attention to the notification about them.*"

A minority of the participants preferred reviewing all the notifications they received carefully and individually. Thus, they preferred placing "reminder" notifications at the bottom of the drawer, because they wanted to handle these "important" to-dos last, i.e., only after they had reviewed and handled all the rest of the notifications. As P20 put it, "*I wanted to put them on the bottom. I didn't want to deal with them right away, but that [placement] would remind me that I still had something to deal with.*" Similarly, P15 commented: "*I place all the ads above my messages so that I can force myself to open and remove all of them before handling my IM messages.*" P32 mentioned wanting these "reminder" notifications placed in the lowest position because it was closest to the "clear all" button. Doing so made him aware of what things remained to be done. "*I hope that after checking all of them and before I press 'Clear all', the last notification I saw would be an important one.*" This also help explain our finding in 4.2.2 that users do not always want to put attractive and important content in higher positions.

5.1.2 Glanceability and Visibility of the Notification. The interviewees also tended to associate placement on the top with high visibility and easy glanceability. They sometimes placed notifications at the top to make it easier to read their titles at a glance. As P04 noted, "*When I am working, I want the news to be on the top. I can't click it but putting those notifications on the top makes it easier to skim their titles*". Others put notifications that were glanceable on top so that they could get the gist of them even when scanning quickly: "*I want the notifications that I can easily comprehend to be placed on the top, because you can know whether you need to handle them.*" (P20).

Several participants mentioned that, given the high visibility of notifications on the top, some types of notifications appeared there too often. As P14 said: "*The weather app appeared too frequently. If it kept showing on the top, it'd be annoying. I want it to be below other notifications.*" P08 also commented about his infrequent interactions with the Major Ongoing class of notifications: "*I wish the music play notifications could be placed last. Because I hope the notifications at the top will be more important. I rarely need to control the music player. I only need to know where it is when I need to control it.*"

The topmost positions were widely considered suitable for notifications that appear only briefly before disappearing of their own accord, "*Maps, Google [Weather], Clock, can be placed on the top, because they usually end pretty fast, and they also less frequently appear*" (P30). A few participants expressed concerns about privacy when certain notifications were highly visible. P13, for example, stated: "*There were other passengers next to me, and I didn't want them to know any of my social media notifications.*"

5.1.3 Positioning as a Means of Clustering. Several participants mentioned the idea of clustering all notifications of a particular kind together into blocks that shared similar interest values. For example, both P10 and P31 wanted to group finance-related notifications together, "*I watch stock after work notifications. Everything I put on the top is about the money*" (P31). Notifications from the same app may have different characteristics, causing them to be clustered into different blocks. As P34 noted, "*If it's purchase-related information instead of Ad information from Shoppee, it should go to the second position below Gmail. Ad notifications should be clustered together no matter which apps they come from.*" The perceived benefits of clustering notifications included not only viewing them together, but also to apply batch operations to them, as P15 explained: "*I wish all ads were in one block. So that I could remove the entire block!*"

5.1.4 Not Needing to Be Attended to, but Needing to Be Displayed. Participants consistently mentioned their feelings toward self-explanatory notifications that they deemed unnecessary to attend to but desired to display in the notification drawer. The major reason was that no further information was needed. As P23 said: "*There are many notifications where the title already tells you enough information. You don't need to tap into them.*" Participants also mentioned that some notifications did not take users to anywhere when acted upon. "*You don't click into temperature [notifications], because there's nothing in them. But they're important, so I ranked these higher in my desired display order than in my click order*" (P32). Other participants noted that they simply preferred to handle certain classes of notification on another device: e.g., "*I respond to Gmail using my computer*" (P15).

5.2 Context Effects

The influence of context, too, was reflected in our interviewees' comments on how their desired display orders often changed. The major types of such influences were location, activity, and temporal contexts: "*If you transferred money on that day, you'd look at that. If it's near some holidays like Father's Day, you'd like to see shopping ads*" (P18). Similarly, P11 said, "*When I was renting a scooter, there* *was the countdown, and you need to monitor its countdown when you're riding it. During this time you want to see it easily*". And P13 noted: "*When I was running in the park, my Samsung Health notifications needed to be shown at a higher position, so that I could click on them easier.*". In short, momentary needs for certain types of notifications to be shown at the top of the drawer were another source of mismatches between what the system displayed and our participants' desired display order.

6 DISCUSSION

This study was motivated by prior research that found users do not always strictly follow a top-down reading order and by perceiving that this might reveal an important limitation – perhaps amounting to a flaw – in the current notification-sorting mechanism; or, it could be attributable to smartphone users' tendency toward selective attendance. We found that two phone-assigned types of notifications (i.e., *Major Ongoing* and *People-to-People*) were assumed by the target phone system to be of high importance, but our participants' often wanted such notifications placed in low positions, or even not shown at all. They also often placed other notifications above these ones.

Furthermore, through our seven-category typology of notifications, we identified three types of mismatches, each of whose attributes differed not only from one another, but all four of the other notification categories. Based on their characteristics, we can conclude that none of these mismatches necessarily reveal notification systems' problems, but rather, strong inter-user variation in notifications' preferred or perceived roles. In the case of *Higher Handled Later*, for example, few notifications were associated with People-to-People communication, but our participants desired them to appear at or near the top of the stack anyway. This was because these notifications' roles were more like reminders or to-do-list items, which made the participants want to notice them readily. *Lower Handled Quicker* notifications, on the other hand, were mostly from news providers, chiefly social apps, shopping apps, and new apps, and attended to quickly because the content was of particular interest to the recipient, and/or interest relatively time-sensitive. Yet, in the grand scheme of things, they were not deemed important enough to be displayed at the top of the drawer. Lastly, *Displayed Not Handled* notifications were more like utility suppliers: not urgent, not attractive, and requiring no particular action, and yet supplying something – in this case, information – that was intermittently required and thus important enough to remain in the notification drawer (sometimes, even at the top) to ensure easy access.

Nevertheless, despite the various roles these notifications played from users' perspectives, the one-dimensional and top-down character of the current notification layout provides only one degree of freedom in how Android phone systems position notifications, and how users attend to them. As a result, it makes sense that while users can only perform simple actions on the current vertical notification layout, with its limited affordances and low control over additional operations, without being able to organizing them as they desire, it is only natural to observe discrepancies between how users want to display their notifications and how the current sorting mechanism actually displays them.

Attendance behavior within the existing one-dimensional space is driven by a mixture of intentions. For instance, not clicking on a notification could be due to lack of interest, but equally, due to the perceived importance of handling it later, or even due to its role as a utility supplier: fulfilling relatively important needs. Likewise, dismissing a notification may be due to lack of interest, but also to an intention to "clear space" for a more important one to remain at (or return to) the top of the drawer. But unfortunately, the inherent limitations of the current one-dimensional notification layout make it challenging at best to infer these intentions.

Other factors that could introduce variance into how users want to display notifications should also be highlighted here. First, the desired display order is also situated; i.e., in different location and activity contexts, our participants expressed different desired display orders. They also mentioned other contextual reasons such as social context (avoiding a companion seeing one's own notifications), temporal context (time-sensitive discounts, seasonal promotions), and momentary needs (notifications related to immediate actions such as a payment transaction). Second, we learned that individual differences affect how people handle notifications (e.g., dismiss all vs. careful review of individual notifications), at least in terms of where they placed the ones they deemed to be "reminders".

Given these observations and system limits, we would not choose to claim that the current notification-sorting mechanism is flawed; nor would we claim, even if that mechanism stands in need of improvement, that the large discrepancies we observed were caused by it alone. Rather, we would argue that the current sorting mechanism makes a reasonable assumption that IM should be prioritized at the top, and that in-progress processes require some attention. These assumptions themselves appear to be well justified, as IM notifications have been found to be favored by smartphone users in numerous studies (e.g. [16, 20, 21]), and easy-entry access and control of ongoing processes also provides good usability [22]. However, we should not ascribe the large discrepancies in the two orders in our data to only the large number of notifications that many people are bombarded with. Instead, it reflects the phones' general lack of consideration of recipients' contexts; current interfaces' lack of affordances for specific operations to be performed on notifications; general lack of awareness of how users attach meanings to particular positions; general lack of consideration of the differential roles of different types of notifications; and finally – very importantly – the limits of one-dimensional layout for accommodating the various roles of notifications from users' perspectives.

With the passage of time, notifications have acquired a rich variety of roles, and many can now be characterized as multi-purpose. Making multi-purpose notifications fit into a one-dimensional layout is the essential root cause of mismatches, and also influences user experience in other ways. If our goal is to make users more efficient at attending to notifications, and to optimize the value of the roles that users want certain notifications to serve, it is time to think about the design space of a two-dimensional notification layout, with additional facilities that should be offered to support the various roles of notifications. Below, we provide our design suggestions.

6.1 Design Suggestions

We offer a series of design recommendations, which can broadly be categorized into short-term and long-term ones. First, we suggest that current notification systems be modified to allow their users to "pin" notifications in specific places, as any given location within the notification drawer may have different meanings to different users. Second, we suggest that current interfaces be adjusted to allow users to manually place their notifications in any order they like. As well as enabling their fulfillment of users' own needs, this will provide opportunities for practitioners to observe users' behavior and preferences, and thus to create more personalized automatic notification-sorting systems. It will also allow practitioners to observe which notifications and notification types win and lose the "duels" between them, thereby learning both personal and general trends in notification-arrangement preferences. And third, we suggest that existing systems be upgraded to include context awareness, since context matters to both attendance and display orders. Such systems could learn from users' sorting results in various contexts, and adjust notification-display order accordingly.

Additional features that can further facilitate the various roles of notifications include 1) highlighting, 2) shortcuts, 3) clustering, and 4) filtering/sorting. Specifically, highlighting makes certain notifications more salient, making it easier for users to notice them. Decisions about whether to apply highlights can be determined by user customization or by the system learning from users' preferences. However, considering that privacy concern is another factor, user permission may be necessary, at least before the first automated highlight is displayed [6]. Shortcuts serve a similar purpose: allowing users to quickly see notifications from specific categories/apps. Clustering, meanwhile, allow notifications from different apps to be reviewed together, as well as batch operations applied to them. And lastly, filtering/sorting features will allow users to quickly find targets and narrow down the notifications they are seeing to specific types.

Yet, while the above features may suffice to leverage the multi-faceted roles smartphone users are already assigning to their notifications, we argue that a better long-term solution would be to provide two-dimensional layout that could accommodate all four of the features described above, while at the same time accommodate more complex information architecture to enable better organization of heterogeneous and multi-purpose notifications. For example, an interface that allows tabbed pages that expand horizontally would make it easier to navigate through different kinds of notifications. As different users are interested in different kinds of digital content (e.g., news, ads, shopping items), customization of tabs and labels should also be provided.

6.2 Research Limitations

One limitation of this study is that we did not compare the original display order and the actual attendance order of the notifications that were simultaneously present in the notification drawer against the participants' desired display order of that list. Though having done this would no doubt have yielded evidence about the discrepancies between the order of notifications in the existing Android system and participants' desired display order and provided more insight into users' actual behavior toward notifications, it would have lowered the diversity of the sampled notifications, meaning that we would not have been able to answer our more important research questions. Since the position in which a notification is placed can be profoundly affected by the characteristics of the other five notifications in the same sorting list, we deemed it important to diversify the combinations of notification categories in each list by using the sampling criterion mentioned earlier, which ensured some level of diversity in every list. Second, we asked the participants to sort the notifications in a flat list despite that current Android system provides grouping the notifications from the same app into a stack[3]. It was because that it would contract our goal to diversify the sampled notifications. Third, participants could only sort six notifications. We were also reluctant to ask our participants to sort more notifications, since in ESM, they had to answer questions about four attributes of every notification they were shown; thus, each additional notification in an ESM list would have added to the number of questions they had to answer. Additionally, we would not necessarily have learned as much from a longer list about how users would change their desired display orders, because a perception that they needed to scroll to see certain notifications might have impacted where they wanted to display them. Fourth, due to our ESM mechanism's design, the participants had to recall their context and perceptions toward notifications 10 minutes prior to when they received ESM questionnaires. Although participants in our pilot study said that it was not hard to recall the moments, our collected data might still suffer from their recall biases. Fifth, there may be more reasons for which users would not want to place important notifications on the top that we did not learn from participants (e.g. notifications at the bottom of the drawer can be reached easily). Sixth, the number of participants was relatively small, young-skewing, and all Android users based on Taiwan. Thus, their behaviors towards and perceptions of phone notifications may not be generalizable to older people, users of other types of smartphones, or residents of other parts of the world. Seventh, it should be noted that, while we captured notification contents and metadata, and phone-usage logs, these were not included in our analysis. Finally, our research and design suggestions are focused on smartphone notifications, further research can be conducted in multi-device environment.

7 CONCLUSION

Using mixed methods to investigate 34 Android users' desired notification-display order, we first revealed a discrepancy between such order and their phones' current display mechanism. Second, we classified seven types of notifications, which allowed us not only to observe the distinctive attributes of each of the seven types of notification, but also to understand the implied roles of each type, and the implied meaning of display orders. This was especially manifested in our identification of the three main species of mismatches between the two orders. In addition, we observed that the multi-purpose roles users tended to assign to their notifications could not be sufficiently supported by their existing notification interfaces, characterized by strict vertical stacking; and for that reason, they desired more freedom to sort notifications on their

[3]Material Design Android notifications: https://material.io/design/platform-guidance/android-notifications.html#behavior

own. Our qualitative findings have revealed some of the variety of the meanings that people tend to attach to particular positions in the notification drawer when arranging their notifications, and provide explanations for the mismatches between display order and attendance order that we captured. Based on these results, we offer several short-term and long-term design recommendations for future notification systems.

ACKNOWLEDGMENTS

This research was supported in part by the Ministry of Science and Technology, R.O.C (MOST 107-2218-E-009-030- MY3)

REFERENCES

[1] Piotr D. Adamczyk and Brian P. Bailey. 2004. 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* (Vienna, Austria) *(CHI '04)*. Association for Computing Machinery, New York, NY, USA, 271–278. https://doi.org/10.1145/985692.985727

[2] Daniel Avrahami, Susan R. Fussell, and Scott E. Hudson. 2008. IM Waiting: Timing and Responsiveness in Semi-Synchronous Communication. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work* (San Diego, CA, USA) *(CSCW '08)*. Association for Computing Machinery, New York, NY, USA, 285–294. https://doi.org/10.1145/1460563.1460610

[3] Jiang Bian, Anlei Dong, Xiaofeng He, Srihari Reddy, and Yi Chang. 2013. User Action Interpretation for Online Content Optimization. *IEEE Trans. on Knowl. and Data Eng.* 25, 9 (Sept. 2013), 2161–2174. https://doi.org/10.1109/TKDE.2012.130

[4] Yung-Ju Chang, Yi-Ju Chung, and Yi-Hao Shih. 2019. I Think It's Her: Investigating Smartphone Users' Speculation about Phone Notifications and Its Influence on Attendance. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services* (Taipei, Taiwan) *(MobileHCI '19)*. Association for Computing Machinery, New York, NY, USA, Article 14, 13 pages. https://doi.org/10.1145/3338286.3340125

[5] Yung-Ju Chang and John C. Tang. 2015. Investigating Mobile Users' Ringer Mode Usage and Attentiveness and Responsiveness to Communication. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Copenhagen, Denmark) *(MobileHCI '15)*. Association for Computing Machinery, New York, NY, USA, 6–15. https://doi.org/10.1145/2785830.2785852

[6] Yi-Shyuan Chiang, Ruei-Che Chang, Yi-Lin Chuang, Shih-Ya Chou, Hao-Ping Lee, I-Ju Lin, Jian-Hua Jiang Chen, and Yung-Ju Chang. 2020. Exploring the Design Space of User-System Communication for Smart-Home Routine Assistants. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376501

[7] Android Developers. 2017. *Notifications UX: What's New for Android O (Google I/O '17)*. Google I/O. https://www.youtube.com/watch?v=vwZi56I0Mi0&t=1170s&ab_channel=AndroidDevelopers [Online; accessed 15-September-2020].

[8] Tilman Dingler and Martin Pielot. 2015. I'll Be There for You: Quantifying Attentiveness towards Mobile Messaging. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Copenhagen, Denmark) *(MobileHCI '15)*. Association for Computing Machinery, New York, NY, USA, 1–5. https://doi.org/10.1145/2785830.2785840

[9] Trinh Minh Tri Do, Jan Blom, and Daniel Gatica-Perez. 2011. Smartphone Usage in the Wild: A Large-Scale Analysis of Applications and Context. In *Proceedings of the 13th International Conference on Multimodal Interfaces* (Alicante, Spain) *(ICMI '11)*. Association for Computing Machinery, New York, NY, USA, 353–360. https://doi.org/10.1145/2070481.2070550

[10] Joel E. Fischer, Nick Yee, Victoria Bellotti, Nathan Good, Steve Benford, and Chris Greenhalgh. 2010. Effects of Content and Time of Delivery on Receptivity to Mobile Interruptions. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services* (Lisbon, Portugal) *(MobileHCI '10)*. Association for Computing Machinery, New York, NY, USA, 103–112. https://doi.org/10.1145/1851600.1851620

[11] Sandy Gould, Duncan Brumby, and Anna Cox. 2013. What does it mean for an interruption to be relevant? An investigation of relevance as a memory effect. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 57 (09 2013), 149–153. https://doi.org/10.1177/1541931213571034

[12] iOS Developers. 2020. *Human Interface Guidelines*. Apple. https://developer.apple.com/design/human-interface-guidelines/ios/overview/mac-catalyst/ [Online; accessed 15-September-2020].

[13] Kostadin Kushlev, Jason Proulx, and Elizabeth W. Dunn. 2016. "Silence Your Phones": Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1011–1020. https://doi.org/10.1145/2858036.2858359

[14] Hao-Ping Lee, Kuan-Yin Chen, Chih-Heng Lin, Chia-Yu Chen, Yu-Lin Chung, Yung-Ju Chang, and Chien-Ru Sun. 2019. Does <i>Who</i> Matter? Studying the Impact of Relationship Characteristics on Receptivity to Mobile IM Messages. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300756

[15] Luis Leiva, Matthias Böhmer, Sven Gehring, and Antonio Krüger. 2012. Back to the App: The Costs of Mobile Application Interruptions. In *Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services* (San Francisco, California, USA) *(MobileHCI '12)*. Association for Computing Machinery, New York, NY, USA, 291–294. https://doi.org/10.1145/2371574.2371617

[16] Hugo Lopez-Tovar, Andreas Charalambous, and John Dowell. 2015. Managing Smartphone Interruptions through Adaptive Modes and Modulation of Notifications. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) *(IUI '15)*. Association for Computing Machinery, New York, NY, USA, 296–299. https://doi.org/10.1145/2678025.2701390

[17] Andrés Lucero. 2015. Using Affinity Diagrams to Evaluate Interactive Prototypes. In *15th Human-Computer Interaction (INTERACT)*, Vol. 9297. Springer, Bamberg, Germany, 231–248. https://doi.org/10.1007/978-3-319-22668-2_19

[18] Afra Mashhadi, Akhil Mathur, and Fahim Kawsar. 2014. The Myth of Subtle Notifications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (Seattle, Washington) *(UbiComp '14 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 111–114. https://doi.org/10.1145/2638728.2638759

[19] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2016. PrefMiner: Mining User's Preferences for Intelligent Mobile Notification Management. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Heidelberg, Germany) *(UbiComp '16)*. Association for Computing Machinery, New York, NY, USA, 1223–1234. https://doi.org/10.1145/2971648.2971747

[20] Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing Content-Driven Intelligent Notification Mechanisms for Mobile Applications. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan) *(UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 813–824. https://doi.org/10.1145/2750858.2807544

[21] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My Phone and Me: Understanding People's Receptivity to Mobile Notifications. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1021–1032. https://doi.org/10.1145/2858036.2858566

[22] Jakob Nielsen. 1994. Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 152–158. https://doi.org/10.1145/191666.191729

[23] Lauren Norrie and Roderick Murray-Smith. 2015. Impact of Smartphone Notification Display Choice in a Typing Task. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (Copenhagen, Denmark) *(MobileHCI '15)*. Association for Computing Machinery, New York, NY, USA, 1094–1099. https://doi.org/10.1145/2786567.2794335

[24] Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Seattle, Washington) *(UbiComp '14)*. Association for Computing Machinery, New York, NY, USA, 897–908. https://doi.org/10.1145/2632048.2632062

[25] Veljko Pejovic, Mirco Musolesi, and Abhinav Mehrotra. 2015. Investigating The Role of Task Engagement in Mobile Interruptibility. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (Copenhagen, Denmark) *(MobileHCI '15)*. Association for Computing Machinery, New York, NY, USA, 1100–1105. https://doi.org/10.1145/2786567.2794336

[26] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-Situ Study of Mobile Phone Notifications. In *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services* (Toronto, ON, Canada) *(MobileHCI '14)*. Association for Computing Machinery, New York, NY, USA, 233–242. https://doi.org/10.1145/2628363.2628364

[27] Martin Pielot, Tilman Dingler, Jose San Pedro, and Nuria Oliver. 2015. When Attention is Not Scarce - Detecting Boredom from Mobile Phone Usage. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan) *(UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 825–836. https://doi.org/10.1145/2750858.2804252

[28] Martin Pielot and Luz Rello. 2015. The Do Not Disturb Challenge: A Day Without Notifications. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI

EA '15). Association for Computing Machinery, New York, NY, USA, 1761–1766. https://doi.org/10.1145/2702613.2732704

[29] Martin Pielot, Amalia Vradi, and Souneil Park. 2018. Dismissed! A Detailed Exploration of How Mobile Phone Users Handle Push Notifications. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Barcelona, Spain) *(MobileHCI '18)*. Association for Computing Machinery, New York, NY, USA, Article 3, 11 pages. https://doi.org/10.1145/3229434.3229445

[30] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-Scalae Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 3055–3064. https://doi.org/10.1145/2556288.2557189

[31] Angela Sasse, Chris Johnson (editors, and Daniel C. Mcfarlane. 1999. Coordinating the Interruption of People in Human-Computer Interaction.

[32] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2019. The influence of concurrent mobile notifications on individual responses. *International Journal of Human-Computer Studies* 132 (2019), 70 – 80. https://doi.org/10.1016/j.ijhcs.2019.07.011

[33] Aku Visuri, N. V. Berkel, T. Okoshi, J. Gonçalves, and V. Kostakos. 2019. Understanding smartphone notifications' user interactions and content importance. *Int. J. Hum. Comput. Stud.* 128 (2019), 72–85. https://doi.org/10.1016/j.ijhcs.2019.03.001

[34] Dominik Weber, Alexandra Voit, and Niels Henze. 2019. Clear All: A Large-Scale Observational Study on Mobile Notification Drawers. In *Proceedings of Mensch Und Computer 2019* (Hamburg, Germany) *(MuC'19)*. Association for Computing Machinery, New York, NY, USA, 361–372. https://doi.org/10.1145/3340764.3340765

[35] Tilo Westermann, Sebastian Möller, and Ina Wechsung. 2015. Assessing the Relationship between Technical Affinity, Stress and Notifications on Smartphones. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (Copenhagen, Denmark) *(MobileHCI '15)*. Association for Computing Machinery, New York, NY, USA, 652–659. https://doi.org/10.1145/2786567.2793684

[36] Sha Zhao, Julian Ramos, Jianrong Tao, Ziwen Jiang, Shijian Li, Zhaohui Wu, Gang Pan, and Anind K. Dey. 2016. Discovering Different Kinds of Smartphone Users through Their Application Usage Behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Heidelberg, Germany) *(UbiComp '16)*. Association for Computing Machinery, New York, NY, USA, 498–509. https://doi.org/10.1145/2971648.2971696