Do Users Have Contextual Preferences for Smartphone Power Management?

Sophie Welber
University of Chicago
sjwelber@uchicago.edu

Valerie Zhao
University of Chicago
vzhao@uchicago.edu

Claire Dolin
University of Chicago
cdolin@uchicago.edu

Olivia Morkved
University of Chicago
omorkved@uchicago.edu

Henry Hoffmann
University of Chicago
hankhoffmann@uchicago.edu

Blase Ur
University of Chicago
blase@uchicago.edu

ABSTRACT

Smartphones must balance power and performance. While most smartphones offer a power-saving mode, they typically provide a binary choice between full performance and monolithic performance degradation (e.g., reducing both screen brightness and processing speed) to save power. Could smartphones improve the user experience by automatically degrading only selected features based on the usage context? To gauge whether preferences for power-saving strategies vary by context, we conducted a 304-participant, survey-based experiment. Each participant was assigned a context (e.g., navigation) and degradation level. They viewed a series of side-by-side simulations of one smartphone operating normally in that context and another operating with reduced GPS accuracy, processing speed, or screen brightness. Participants rated their willingness to accept each tradeoff to save power. Contrasting current power-saving modes, we found that participants' preferences did indeed vary by context. Using factor analysis to cluster preferences, we identified key personas that pave the way toward context-aware and self-aware alternatives to smartphone power-saving modes.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; **Smartphones**.

KEYWORDS

human-battery interaction, smartphones, self-aware computing, adaptation, power saving

ACM Reference Format:

Sophie Welber, Valerie Zhao, Claire Dolin, Olivia Morkved, Henry Hoffmann, and Blase Ur. 2021. Do Users Have Contextual Preferences for Smartphone Power Management?. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21), June 21–25, 2021, Utrecht, Netherlands.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3450613.3456813

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '21, June 21–25, 2021, Utrecht, Netherlands

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8366-0/21/06…$15.00
https://doi.org/10.1145/3450613.3456813

1 INTRODUCTION

Smartphone batteries last for a limited time on a single charge, and fears of running out of battery make users anxious [31]. Most smartphones have a power-saving mode that automatically prolongs battery life by sacrificing performance, but these modes typically do not allow users to customize which aspects of performance to degrade [4, 21]. Users instead turn to ad hoc strategies, such as closing background apps. Unfortunately, these strategies are often ineffective [18, 22], especially when users have an incorrect mental model of their phone battery [16, 35]. Manual adjustments can also be time consuming since they require readjustment whenever the user's context or goals change.

These shortcomings could be alleviated with a self-aware system that automatically adjusts battery consumption to meet the user's power and performance goals in a given situation [26]. Researchers have investigated how to implement power management based on end-user goals [20, 29], but surprisingly little work has employed user-centered research to determine what these goals actually are. To the best of our knowledge, most smartphone power-saving systems do not consider context when making power-performance tradeoff decisions. This gap has motivated the Adaptive Battery feature of Android Pie [2, 34], though it only considers the frequency of use of specific apps. Instead, we aim to improve existing power-saving modes by taking into account users' power-performance goals and the factors that affect them.

We hypothesized that user preferences would differ depending on the task a user was doing on their phone (**context**) and on the degree of the performance degradation (quality **level**). In a survey-based, 304-participant online experiment with a mixed-subjects design, we studied willingness to accept three power-saving performance degradations (**tradeoffs**): reduced GPS accuracy, processing speed, and brightness. We randomly assigned each participant to a context and level. For each of the three tradeoffs, we showed them a comparison of two smartphone screens in their assigned context, one running normally and one degraded to their assigned level. We asked participants how often they would accept this tradeoff given a corresponding amount of power savings, comparing participants' responses between the assigned contexts. To simulate these tradeoffs realistically, we based our demonstrations on real measurements of each tradeoff on a Google Pixel. To allow for additional analysis across contexts, we also asked participants about their preferences for all relevant context-tradeoff combinations.

We found that GPS acceptance varied across contexts and levels, and processing speed acceptance varied across contexts. While brightness acceptance did not vary significantly across contexts in the between-subjects portion of our study, it did vary significantly in the within-subjects portion. Participants expressed a variety of concerns, such as inaccurate GPS readings while navigating unfamiliar places (yet not wanting to run out of battery), sensitivity to reduced processing speed when playing games, and reluctance to reduce screen brightness when working with photos. More generally, participants indicated frequent worries about running out of battery on their phone, and they reported often using both automatic power-saving modes and manual adjustments to save power.

To better understand cross-cutting patterns in participants' preferences, we also clustered their responses using factor analysis. As one might expect, we found that participants' overall willingness to accept a given tradeoff was correlated across contexts. In addition, we identified a cluster representing participants reluctant to accept reduced phone brightness and speed while playing a game, and another representing participants reluctant to accept reduced phone speed while taking a photo. We also observed a larger cluster representing participants reluctant to accept any tradeoffs while navigating, as well as reluctant to accept mission-critical tradeoffs while taking a photo or watching a video. Our results pave the way toward better modeling users as part of developing context-aware and self-aware power-saving modes that conserve power while preserving much of the phone's user experience, different from the monolithic approach of current power-saving modes. We discuss considerations for transferring these results to practice.

2 RELATED WORK

We build on existing literature on human-battery interaction (HBI). Rahmati et al. examined user concerns about battery life and misconceptions about batteries, proposing improvements for battery indicators [35]. Various studies have piloted interactive battery interfaces [17, 39] and battery-awareness applications [6] to determine their effect on user behavior. Other works have studied user charging habits [10, 16]. Recent work has studied human-battery interaction in specific populations, such as among cultural groups [14, 32] or users with visual impairments [40]. Hosio et al. prompted study participants to attach a monetary value to battery life, concluding that battery value is affected by context [27]. We contribute to this literature by studying whether user acceptance of power-performance tradeoffs is also contextual.

Outside of HBI, prior work has examined battery performance, establishing technical benchmarks while acknowledging that users differ in their priorities. Lee et al. sought to improve quality of service by consulting users' battery lifetime goals [29]. Other works that focus on technical improvements have asserted that users differ in their tolerance for battery depletion [38] and their energy goals [24]. These works suggest that users should not be left out of battery-management decisions, yet have not systematically studied users' preferences and tradeoff decisions. We fill this gap.

Technological advances have made mobile devices more capable of modifying their own processes as part of context-aware [9, 13, 19, 37] or self-aware [1, 25, 26, 30] computing systems. Prior work on context-aware battery management proposed improvements to mobile systems [11, 23, 28] and to mobile applications [41]. Ravi et al. used context awareness to aid mobile devices in predicting when the device would next be charged [36]. Martins et al. proposed an abstracted interface that lets users make informed power tradeoffs while allowing applications to react autonomously to these choices [33]. Hoffmann et al.'s framework for self-aware systems asserts that "self-aware computational models [should] automatically adjust their behavior in response to environment stimuli to meet user specified goals" [26], motivating us to elicit users' specified goals for mobile battery usage. Our work aids the development of context- and self-aware battery systems by presenting insights into the tradeoffs users would want their phones to make under various conditions. Future work can leverage these insights to develop self-aware power-saving modes that improve battery life without meaningfully impacting the user experience on a phone.

3 METHOD

We conducted an online, survey-based study of smartphone users in August 2020. The survey focused on what performance tradeoffs a user would accept, as well as how these preferences varied across contexts. We validated the survey through pilot testing and cognitive interviews. We recruited participants on Prolific. We required participants be 18+ years old, live in the US, regularly use a smartphone, have a 95%+ approval rating on Prolific, and take the survey in the Firefox, Safari, or Chrome web browser with JavaScript enabled (necessary for our side-by-side videos). We compensated $5 USD for this IRB-approved, 30-minute study.

Conditions. To collect information on tradeoff preferences for a variety of situations, part of the study used a between-subjects design where we randomly assigned participants to one of six **contexts** detailed in Table 1: playing a game, navigating with a map, taking photos, using social media, watching a video, or web browsing. Since there is no baseline way to use a smartphone, we chose to study a sample of contexts based on tradeoff-relevant apps commonly used on Apple and Android phones [3, 7, 8, 12].

We tested three performance **tradeoffs** detailed in Table 2: reduced screen brightness, processing speed, and GPS accuracy. Reduced screen brightness is achieved by lowering the screen brightness setting below the level set by the user or phone. Reduced processing speed is achieved by lowering the Dynamic Voltage and Frequency Scaling (DVFS), which affects the power and speed settings on the phone's processors. Reduced GPS accuracy polls for location less frequently. While this latter tradeoff cannot currently be achieved, even on a rooted phone, because of the proprietary nature of current GPS apps, we simulated the tradeoff with a GPS spoofing app that removes a fraction of the points in a GPS trace and then proportionally raises the polling frequency.

Not all tradeoffs apply to all contexts. We asked all participants about reduced screen brightness and reduced processing speed, but only asked participants in the *navigation*, *game*, and *photo* conditions about reduced GPS accuracy. These three contexts are the only ones where an app would collect GPS data periodically for the app's main function. Navigating with a map relies on updated GPS data on the user's location, some games (e.g., Pokemon Go) use a person's real-time location for game play, and camera apps record location information in photo metadata. While apps in the other contexts can collect location data, it is either not typically part of those apps' main functionality or the location data is collected only

Table 1: The six contexts and how they were described to participants. Each participant was randomly assigned to one context.

Context	Description Presented to Participants
Gaming	When we say "playing a game," we include anything that you consider to be a game. Some examples of apps that people use for playing a game are: Candy Crush, Sudoku, Fruit Ninja.
Navigation	When we say "navigating with GPS," we mean using your phone to direct you somewhere using GPS. This could include reading the directions off the screen or having the phone display and speak the directions out loud. This only includes times where you are actively moving towards your destination, so this does NOT include, for example, looking up directions before you leave your house.
Photo	When we say "taking photos," we mean using your phone as a camera to take a photo. This does NOT include editing or viewing your photos. Some examples of apps that people use for taking photos are: Google Camera, Apple Camera, Snapchat.
Social Media	When we say "using social media," we mean using your phone to look at, post on, or otherwise interact with a social media website. Some examples of apps that people use for using social media are: Facebook, Instagram, Twitter.
Video	When we say "watching a video," we mean using your phone to watch a video from any source. Some examples of apps that people use for watching a video are: YouTube, TikTok, Netflix.
Web Browsing	When we say "web browsing," we mean any activity you do on your phone in a web browser. This probably means going to a website. This does NOT include internet-connected apps that have specific purposes, like a social media or weather app. Some examples of apps that people use for web browsing are: Safari, Chrome, Firefox.

Table 2: The descriptions provided to participants explaining each tradeoff.

Tradeoff	Description
GPS	When GPS accuracy is normal, your phone will sync frequently with GPS satellites, so your device will have accurate and frequently updated location information. When GPS accuracy is reduced, your phone will sync less frequently, so it may not know your exact location.
Speed	When processing speed is normal and you tap on your phone screen, the device will respond immediately. When processing speed is reduced, you may notice a lag between your actions and your phone's response.
Brightness	When screen brightness is normal, your screen should be easy to see. When screen brightness is reduced, your screen may look dimmer than you would normally want it to be.

Table 3: Each participant was assigned a quality level indicating the amount of performance degradation and the accompanying increase in time the phone would last on a charge, as compared to lasting for one hour. This table lists those mappings.

Tradeoff	Description	Low	Medium	High
GPS	Reduced GPS accuracy	1 hour, 31 minutes (1.51 hours)	1 hour, 21 minutes (1.35 hours)	1 hour, 14 minutes (1.24 hours)
Speed	Reduced processing speed	2 hours, 7 minutes (2.12 hours)	1 hour, 54 minutes (1.90 hours)	1 hour, 31 minutes (1.51 hours)
Brightness	Reduced screen brightness	1 hour, 14 minutes (1.23 hours)	1 hour, 13 minutes (1.22 hours)	1 hour, 12 minutes (1.21 hours)

once (e.g., to display the local weather) and is thus unaffected by polling frequency.

To understand how user preference varies with the degradation magnitude and accompanying battery savings, we also randomly assigned participants to one of three quality **levels**: low, medium, or high. In the high condition, the degradation and subsequent power savings are small. In the low condition, the degradation and subsequent power savings are larger. To determine how much of the degradation to apply at each level, we collected power usage data from an automated web-browsing benchmark on a rooted Google Pixel phone. We chose three locations on these power usage curves that represented sufficiently distinct power levels and used the corresponding level of reduction. Table 3 presents the battery savings for each level, which was shown to participants to inform their decision about the tradeoff.

Survey Instrument. After completing a consent form, participants indicated how often they use their phone for their assigned context and provided examples of apps they use in that context. We then elicited their preferences regarding the two or three tradeoffs relevant to their assigned context in randomized order in the **between-subjects portion** of our study. This portion of the study aimed to unpack individual preferences. For each tradeoff, we provided a text description and side-by-side visual example contrasting what an app in that context would look like normally and with the tradeoff applied, as shown in Figure 1. We also noted how much extra time the battery would last with the tradeoff (e.g., "for every 1 hour of use, the phone would instead last for 1.5 hours"). We required participants to watch the side-by-side videos for their complete duration. We then asked participants to indicate if they would *always*, *sometimes*, or *never* accept the given tradeoff at the given level. If a participant answered *sometimes*, we asked them to describe the situations in which they would and would not accept the tradeoff. For participants who answered *always* or *never*, we asked them to explain their reasoning and describe situations where they would have answered differently.

While the between-subjects portion gave us rich data about specific preferences and underlying rationales, we also wanted to understand how preferences for different tradeoffs and contexts broadly related to each other. Thus, we also conducted a **within-subjects portion**. We asked participants to rate, on a 5-point Likert scale, how likely they would accept *each* of the three tradeoffs for *each* of the six contexts, without additional visual examples or notes about battery savings. These questions thus represent 18 **tradeoff-context pairs**. Recall that participants watched videos

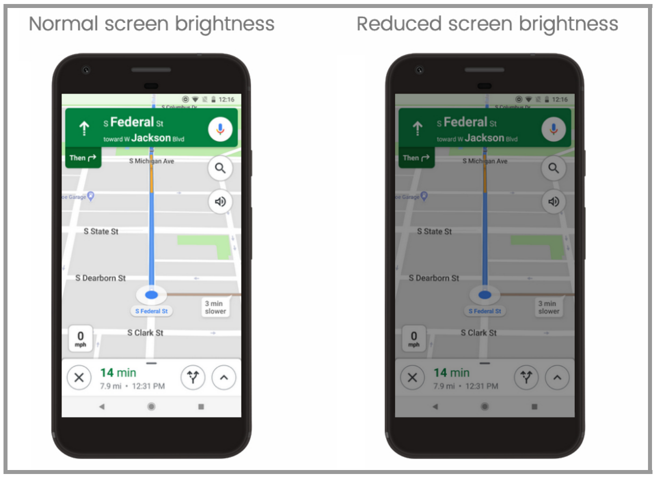

Figure 1: An example of the primary comparison task. Participants saw a pair of videos or images demonstrating normal phone behavior (left) in the assigned context alongside a synchronized version demonstrating a tradeoff (right) at the assigned quality level.

in the between-subjects portion. Half of the contexts were not relevant to GPS, so they did not see videos about GPS. In piloting, we found that asking those participants about GPS tradeoffs was confusing. Thus, half of the participants responded only to speed and brightness tradeoffs in the within-subjects portion, resulting in only 12 (not 18) tradeoff-context pairs.

We also elicited participants' smartphone habits and demographics. We asked participants a series of questions about their general smartphone usage, interactions with power-saving modes, and charging behaviors. Finally, we collected participants' technical expertise, phone purchasing preferences, and basic demographics.

For the tradeoffs, we based our visual examples and estimates of power savings on real-world measurements. We created the examples using the most popular free app related to each task that was available on the Google Pixel phone. The images for the screen brightness tradeoff came from phone screenshots and photo-editing software. Videos of reduced GPS accuracy were screen recordings taken directly from the phone. For processing speed, screen recordings did not accurately communicate the performance degradation, so we simulated it by recording on a phone with normal processing speed and then reducing the frame rate of the video to match the level of reduction users would have experienced. To determine the extra time saved by each tradeoff for each level, we collected power usage benchmark data to calculate the percentage change in power usage between the reduced and non-reduced trials. We then converted the percentage change into a time savings relative to the baseline (selected via pilot testing) of one hour.

Analysis. Quantitative analyses followed three templates. First, we aimed to understand how participants' willingness to accept each of the three tradeoffs varied by context and level, as well as the participant's demographics. We thus built an ordinal logistic regression model for each tradeoff. The willingness to accept the tradeoff was the dependent variable. We included the following independent variables, or *IVs*: the assigned context (categorical), the assigned quality level (ordinal), the interaction between the context and level, the participant's age range (ordinal), their gender (categorical), their technical background (categorical), how long their phone battery lasts (continuous), how frequently they worry about running out of battery (ordinal), and how frequently they use their phone for the assigned context (ordinal). We used backward-elimination by AIC to build a parsimonious model. We focus on the odds ratio, which is the change in odds for a unit increase (ordinal and continuous IVs) or relative to the baseline (categorical IVs).

Our second set of quantitative analyses sought to understand how participants' responses varied across related questions (e.g., comparing a participant's willingness to accept a GPS tradeoff across each of the six contexts). Because these data are not independent, we used the Friedman test, a non-parametric analogue of the repeated measures ANOVA. We first performed the Friedman test across all groups of interest for a particular research question. If this omnibus test was significant, we then performed post-hoc, pairwise comparisons of groups using Eisinga et al.'s method for pairwise comparisons of Friedman rank sums [15]. To control for type I errors in these pairwise comparisons, we used Holm correction and report the corrected p-values. For all of our statistical testing, $\alpha = .05$.

Third, to identify clusters of preferences among participants, we performed factor analysis. Specifically, we used maximum-likelihood factor analysis with a varimax rotation across the 18 tradeoff-condition pairs in the within-subjects portion of our study. Because half of the participants were assigned a context for which GPS tradeoffs did not apply, we included in our first model only the half of participants who had been asked all 18 questions. We repeated this process for the 12 tradeoff-condition pairs for speed and brightness, which all participants answered.

We performed qualitative analysis of answers to free-response questions regarding decisions to *always*, *sometimes*, or *never* accept a tradeoff. With three tradeoffs to ask about, three answer choices, and two free-responses questions asked per choice, we looked at data from 18 questions. For each question, one coder created a codebook through open coding. A second coder independently used that codebook to code all data, and the coders met to resolve discrepancies. Krippendorff's α (reliability) was 0.79.

Limitations. Our study could not fully immerse participants in experiencing each tradeoff in their daily life due to constraints of the COVID-19 pandemic and survey-based study designs. We explicitly instructed participants to report their habits from before COVID-19. Because many participants were working from home during the time of the study, some participants may not have been able to remember their pre-COVID habits clearly.

We showed changes to a smartphone through a simulation observed on a computer screen. Participants relied on text descriptions and side-by-side images or videos to understand how the tradeoffs would affect their phones, but did not experience the changes on their own devices while performing their usual activities. The side-by-side comparison was necessary to provide a baseline for the simulation, but this presents a limitation as users may have reacted more strongly to changes when seeing them alongside the original configuration. With reduced GPS accuracy, participants saw how the apps would react, but did not experience the tradeoff while actually navigating. With reduced screen brightness, participants' perceptions may have been affected by their computers' screen brightness or ambient lighting. For the *photo* context specifically, participants saw a slideshow of screenshots of the map data attached to an image. These examples may have been more difficult to understand if participants were unfamiliar with photo metadata. For the speed tradeoff, participants were not tapping and scrolling themselves, so they may not have been able to fully judge the effect of the tradeoff on the phone's responsiveness. Finally, the brightness examples could have looked different to participants with different screen-brightness settings on their computers.

Despite these limitations, our simulations provide important data. Users' contextual preferences, or lack thereof, regarding smartphone performance-power tradeoffs have not been widely studied. These simulations provide a highly controlled way to elicit user perceptions with consistent experiences for participants. Future work could communicate these tradeoffs by modifying the user experience on participants' own phones, using our findings from these simulations to set initial quality levels and identify contexts likely to have a large impact on preferences.

4 RESULTS

First, we summarize our participants' demographics and phone-use habits. Then, we examine the within-subjects data to compare acceptance of tradeoffs across all six contexts, employing factor analysis to identify personas that predict users' responses for a grouping of tradeoff-context pairs. Next, we examine the between-subjects data to determine additional factors that affected participants' acceptance of each tradeoff. We discuss the qualitative factors that participants frequently mentioned for all three tradeoffs, and then we present a regression model and additional reasoning for each tradeoff individually. Finally, we examine participants' relative preferences about a tradeoff within one app context.

4.1 Participant Demographics and Phone Usage

A total of 350 crowdworkers completed the survey, which took a median of 22.9 minutes and a mean of 27.7 minutes. We excluded 46 participants for failing attention or quality checks, leaving us with 304 participants. Among participants, 55.9% were men, 41.8% were women, and 2.3% were non-binary. Participants were younger than the general population: 22.7% were 18–24 years old, 41.8% were 25–34 years old, 24.3% were 35–44 years old, and the remaining 11.2% were age 45+. Overall, 32.9% of participants reported having a technical background (defined as a degree or job in CS, IT, or a similar field), while the remaining 67.1% did not.

The majority of participants (73.7%) owned only one smartphone, while 24.7% owned two and 1.6% owned three or more. Apple phones were the most popular among participants, with 44.7% reporting Apple as their primary phone's brand. In addition, 30.9% reported that their primary phone was made by Samsung, and 4.9% by Google. Participants reported depending heavily on their phone. Among participants, 34.9% reported using their primary phone at least four hours a day, and 95.4% reported using it for at least an hour a day. Notably, 84.6% of participants "somewhat" or "strongly" agreed that it would be a problem if their primary smartphone "ran out of battery early in the day and [they] couldn't charge it for the rest of the day." Participants rated the importance of 13 characteristics when deciding to buy a phone on a five-point Likert scale ("not at all important" to "extremely important"). The two highest rated categories were phone performance (88% "very" or "extremely" important) and battery life (87%), our key tradeoff.

Running out of battery was a major concern. Participants reported that, on a single charge, their primary phone would last an average of 10.8 hours (1Q: 6 hours; median: 8 hours; 3Q: 12 hours). Critically, 38.2% of participants worried about their battery dying at least once a day, while 64.8% worried about it at least once a week. In contrast, only 12.5% of participants never worried. Participants reported wide use of phone power-saving modes. Overall, 81.6% of participants used such a mode, with 39.8% reporting that their phone enables it automatically when the battery is low and 41.8% reporting that they turn it on manually. In contrast, only 17.1% of participants reported not using such a mode.

4.2 Expected Willingness to Make Tradeoffs

As previously mentioned, we collected participants' preferences both within-subjects and between-subjects. We use the within-subjects data, presented in this sub-section and the following sub-section, to understand relationships between preferences, as well as the relative importance of tradeoffs. We later use the between-subjects data to better understand individual preferences.

In the within-subjects portion, we asked all participants to rate their likelihood to accept *each* tradeoff in *each* of the six contexts, leading to 18 tradeoff-context pairs. For example, "How likely would you be to accept reduced [*screen brightness*], like you saw in the examples on the previous page, in order to save battery when using your phone for [*navigating with a map*]?" Answers were on a five-point Likert scale from "extremely likely" to "extremely unlikely." Recall that participants in the between-subjects portion saw videos about an assigned context. Only half of those contexts were relevant to GPS tradeoffs and therefore only half of participants saw a video about GPS. Participants who did not see a video about GPS did not answer within-subjects questions about GPS.

Figure 2 shows the distributions of these preferences for the three tradeoffs. As we hypothesized, preferences were contextual. First, responses for the GPS tradeoff varied by context (Friedman $\chi^2(5) = 114.87$, $p < .001$). Participants were least likely to accept reduced GPS accuracy in the navigation context, with only 23.4% of participants rating themselves somewhat or extremely likely to do so. Participants were more likely to accept the GPS tradeoff in all five other contexts: social media ($p < .001$), video ($p < .001$), web browsing ($p < .001$), photo ($p < .001$), and game ($p < .001$). Between 50.3% and 61.7% of participants rated themselves somewhat or extremely likely to accept the GPS tradeoff in those contexts. No other pairwise differences between contexts were significant.

Responses for the speed tradeoff also varied by context (Friedman $\chi^2(5) = 111.99$, $p < .001$). Participants were more likely to accept reduced speed while using social media than in the navigation ($p < .001$), game ($p < .001$), video ($p < .001$), or photo ($p = .003$) contexts. Whereas 60.8% of participants were somewhat or extremely likely to accept reduced speed when using social media, only between 37.5% and 50.6% answered similarly for those other four contexts.

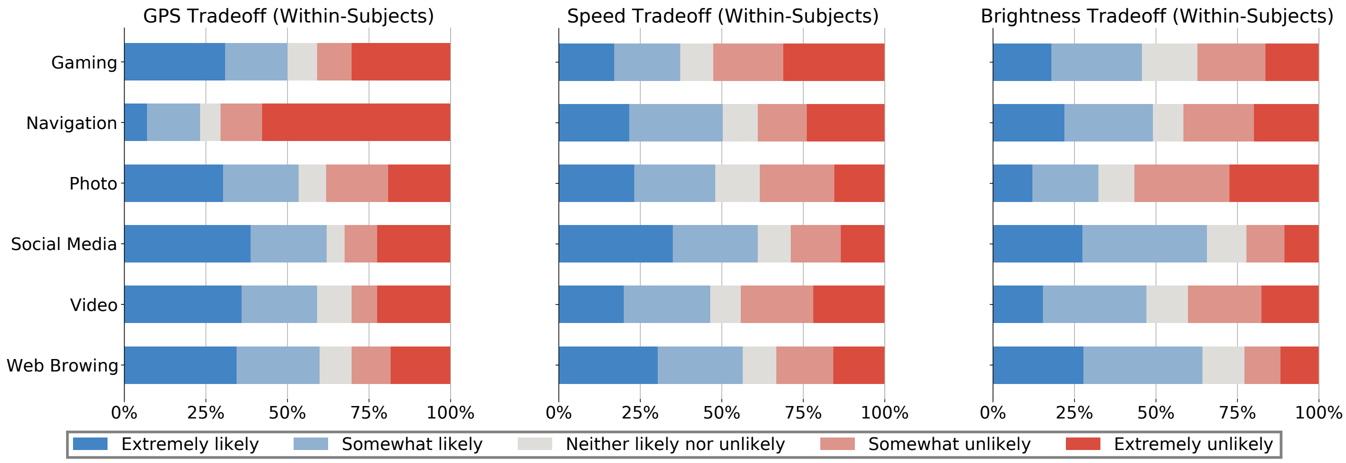

Figure 2: Toward the end of the survey, participants reported their likelihood for accepting their assigned tradeoff in the six contexts (within-subjects). These graphs show participants' responses for these abstract considerations of tradeoffs.

Furthermore, the 56.6% of participants who were somewhat or extremely likely to accept reduced speed while web browsing was significantly more than for navigation ($p = .020$), game ($p < .001$), and video ($p = .002$). Participants were also less likely to accept the speed tradeoff in the game context than in navigation ($p = .010$) and photo ($p < .001$).

Finally, responses for the brightness tradeoff again varied by context (Friedman $\chi^2(5) = 181.14$, $p < .001$). Participants were more likely to accept reduced brightness while web browsing than in the navigation ($p < .001$), game ($p < .001$), video ($p < .001$), or photo ($p < .001$) contexts. Whereas 64.2% of participants were somewhat or extremely likely to accept reduced brightness while browsing the web, this percentage varied from 32.6% to 49.3% for those four other contexts. The 65.5% of participants who were somewhat or extremely likely to accept reduced brightness while using social media was also significantly more than those same four contexts: navigation ($p < .001$), game ($p < .001$), video ($p < .001$), and photo ($p < .001$). Furthermore, the 32.6% of participants who were somewhat or extremely likely to accept reduced brightness in the photo contexts, the lowest of the six contexts, was also significantly less than the navigation ($p < .001$), game ($p < .001$), and video ($p = .002$) contexts.

4.3 Clustering Preferences via Factor Analysis

While preferences are inherently subjective, a power-saving mode can best balance tradeoffs if it can predict these subjective preferences without requiring much input from the user. To that end, we used factor analysis to search for "personas," or clusters in preferences for a participant. If a phone can identify a given user's persona with one question and then accurately predict that user's preferences for a number of different pairs of tradeoffs (e.g., speed) and contexts (e.g., gaming), it can save battery while respecting a user's preferences.

We performed maximum-likelihood factor analysis (using varimax rotation) across the 18 tradeoff-context pairs for the half of participants who had seen videos about the GPS tradeoff. Note that we performed factor analysis for the data from all participants in the 12 tradeoff-context pairs excluding GPS and found highly similar results, so we present only the former. Through factor analysis, we identified a smaller number of latent factors, each of which we call a persona and assign a descriptive name. Each of the original tradeoff-context pairs has a loading (-1 to 1) on each latent factor. Loadings close to 1 indicate that the latent factor captures nearly all of that variable's variance, whereas those close to -1 indicate it captures it with inverse polarity. Numbers close to 0 indicate that the latent factors capture little of the original pair's variance.

To capture the 18 tradeoff-context pairs, using six latent factors was sufficient ($\chi^2(60) = 61.8$, $p = 0.410$), whereas five latent factors was marginal ($\chi^2(73) = 92.3$, $p = 0.063$). Cumulatively, the six factors explained 68.7% of the variance. While this does not explain all of the variation in the data, these factors can still help determine the appropriate contexts for making a power-performance tradeoff.

The first three factors (personas) respectively cluster subjective preferences about the three tradeoffs across contexts. The persona we named *GPS Overall* has large loadings on preferences about the GPS tradeoff in five of the six contexts (all except navigation). The *Speed Overall* persona has large loadings on preferences about the speed tradeoff in all six contexts, and these are the largest loadings on any factor for all of these contexts other than navigation. Similarly, the *Brightness Overall* persona has large loadings on preferences about the brightness tradeoff in all six contexts, and these are the largest loadings on any factor for all of these contexts other than navigation and photo. In short, participants who cared about a particular tradeoff in one context also tended to care about that same tradeoff in other contexts.

The remaining three factors (personas) capture more specific aspects of participants' subjective preferences. First, the *Navigation & Media* persona most strongly captures tradeoffs that are mission-critical during navigation, as well as while taking photos or watching videos. In particular, this persona has strong loadings on all three tradeoffs in the navigation context, representing the heaviest loading on any factor for all three. In addition, this persona also has a heavy loading on the brightness tradeoff in the photo context, as well as weaker loadings on the speed tradeoff in

Table 4: Factor analysis on participants' likelihood to accept speed, brightness, and GPS tradeoffs in the six contexts. The largest loading per tradeoff-context pair is bolded. We omit loadings < 0.25. There were 141 participants (those who saw GPS).

Tradeoff	1: GPS Overall	2: Speed Overall	3: Brightness Overall	4: Navigation & Media	5: Gaming Critical	6: Photography Latency	Communality
Speed in gaming		**0.70**			0.43		0.72
Speed in navigation		0.36		**0.61**			0.57
Speed in photo		**0.59**		0.29		0.51	0.74
Speed in social media		**0.75**					0.65
Speed in video		**0.76**		0.28			0.69
Speed in web browsing		**0.79**					0.75
Brightness in gaming			**0.70**		0.44		0.70
Brightness in navigation			0.40	**0.56**			0.48
Brightness in photo			0.38	**0.57**			0.54
Brightness in social media			**0.88**				0.86
Brightness in video			**0.60**	0.37			0.61
Brightness in web browsing			**0.87**				0.83
GPS in gaming	**0.73**						0.57
GPS in navigation				**0.62**			0.44
GPS in photo	**0.81**						0.71
GPS in social media	**0.90**						0.85
GPS in video	**0.87**						0.80
GPS in web browsing	**0.89**						0.87

the photo context and both the brightness and speed tradeoffs in the video context. In short, this persona captures that some participants had strong preferences about being unlikely to accept battery-saving, performance-degrading tradeoffs particularly in the navigation context, as well as in the tradeoffs most relevant to taking a photo (particularly screen brightness) or watching a video.

Next, the *Gaming Critical* persona captures that some participants had strong preferences against accepting battery-saving, performance-degrading tradeoffs in speed or brightness while playing a game on their phone. In some sense, this persona captures the preferences of "gamers," who care about full performance while playing a game, whereas non-gamers might not care about having degraded performance while playing a game if it saves battery. The final persona, *Photography Latency*, captures only the single tradeoff-context pair of speed while taking a photo. Essentially, some participants expected to be especially sensitive to the phone's speed while taking a photo, whereas others did not. As we discuss further in Section 5, we imagine that a self-aware system would ask the user a few questions or generalize based on some observations of behavior to place a given user into these particular personas, roughly approximating their subjective preferences for context-aware modes for saving battery on a phone.

4.4 Factors Correlated with Acceptance

While the previous analyses reported on our within-subjects data, we now discuss our between-subjects data. Our main response variables were participants' willingness to accept reduced GPS accuracy, processing speed, or screen brightness in their assigned context and at their assigned quality level. Before presenting how preferences differed along these factors, we first describe two factors that participants considered regardless of the tradeoff of interest.

Across all three tradeoffs, an important factor in participants' decisions was their phone's battery status. One of the most common reasons for always accepting a tradeoff was to save battery. For participants who selected *always*, 49/90 (54.4%) mentioned battery life as a reason for accepting the brightness tradeoff, 9/21 (42.9%) for the GPS tradeoff, and 60/132 (45.5%) for the speed tradeoff. Participants who would *sometimes* accept the tradeoff often mentioned battery life as well, with three distinct factors that affected their choice: access to a charger, current battery level, and anticipated battery needs. As an example of the charger access theme, when asked about accepting reduced screen brightness, P53 stated, "If I couldn't charge my phone for a long time I would accept it." Charger access was mentioned in 34/160 (21.3%) brightness responses, 25/115 (21.7%) speed responses, and 2/54 (3.7%) GPS responses. As an example of the current battery level theme, P40 was willing to accept reduced GPS accuracy "if my battery was very low." Low battery was mentioned in 33/160 (20.6%) brightness responses, 10/54 (18.5%) GPS responses, and 14/115 (12.2%) speed responses. Regarding anticipated battery needs, when asked about accepting reduced processing speed, P128 stated, "I can reduce the web browsing speed if I wanna save my battery life for some other tasks I wanna do." Future battery needs featured in 17/160 (10.6%) brightness responses, 6/54 (11.11%) GPS responses, and 15/115 (13.0%) speed responses.

Additionally, many participants mentioned characteristics of navigation routes that affected their decision, such as: familiarity with the area, number of turns, length of drive, and traffic volume. For example, when asked about accepting the brightness tradeoff, P95 stated, "During a dark night, on a particularly straight stretch of road with no turns or other directions - my screen would not need to use full brightness to display information to me." Similarly, P35 was willing to accept reduced GPS accuracy for "low traffic areas, or maybe areas I'm somewhat familiar with," and P54 was okay with reduced processing speed "on long drives or anywhere there aren't quick road changes."

4.4.1 Speed. In our regression (Table 5), we found that participants' preferences about the speed tradeoff varied by context, as further visualized in Figure 3. Compared to the web browsing context, participants were significantly less likely to accept the speed tradeoff in three of the five other contexts: gaming (0.38× as likely, p = .014), navigation (0.44× as likely, p = .023), and photo (0.39× as likely, p = .017). Furthermore, the difference in one of the two remaining contexts, social media (0.51× as likely, p = .068), was marginally significant.

Table 5: Our parsimonious ordinal logistic regression model for respondents' willingness to accept a *speed tradeoff* never (0), sometimes (1), or always (2) based on the app context, the quality level, and the participant's demographics. For categorical IVs, we indicate the baseline category. For ordinal IVs, we indicate the fitted function. We indicate the odds ratio (*OR*), coefficient (*Coeff.*), and standard error (*SE*).

Factor	Baseline	OR	Coeff.	SE	z	p
Context: Gaming	Web Browsing	0.384	-0.957	0.388	-2.466	**0.014**
Context: Navigation	Web Browsing	0.443	-0.815	0.358	-2.279	**0.023**
Context: Photo	Web Browsing	0.390	-0.941	0.395	-2.380	**0.017**
Context: Social Media	Web Browsing	0.505	-0.683	0.375	-1.824	0.068
Context: Video	Web Browsing	0.794	-0.231	0.363	-0.635	0.526

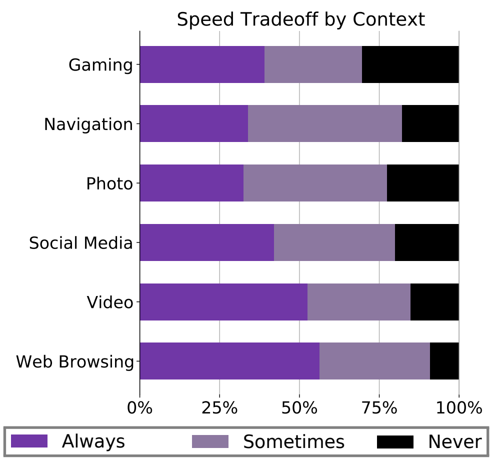

Figure 3: How likelihood to accept a *speed tradeoff* varied across contexts.

The most common reason participants selected for *always* accepting the speed tradeoff was that they felt that the difference between the examples shown was insignificant. For example, "I couldn't tell a difference in the two" (P13) or "the one on the right is still accurate enough" (P14). The second most common reason was to preserve battery life. In contrast, 52.6% of the 57 participants who selected *never* expressed a general preference for normal processing speed or an annoyance with the reduced version, such as "I like things fast and not delayed" (P204).

For the 115 participants who selected *sometimes*, many mentioned the battery or navigation route contexts previously discussed. The other biggest reason for not accepting the tradeoff was the time-sensitivity of the task: "I would not accept reduced processing speed in games where timing was a crucial element. For example, games that included a timer or countdown, or games where speed mattered, I would not accept reduced processing speed" (P7). 22.6% of participants mentioned a time-sensitive task in their response.

4.4.2 Brightness. In our regression (Table 6), we found that only the participant's age was significantly correlated with their preferences about the brightness tradeoff, as further visualized in Figure 4. Older participants were 0.23× as likely as younger participants to accept the brightness tradeoff (p = .016). Even younger participants' free-text justifications sometimes encoded age-related factors, such

Table 6: Our parsimonious ordinal logistic regression model for respondents' willingness to accept a *brightness tradeoff*.

Factor	Baseline	OR	Coeff.	SE	z	p
Age Range	(linear fit)	0.225	-1.492	0.618	-2.417	**0.016**

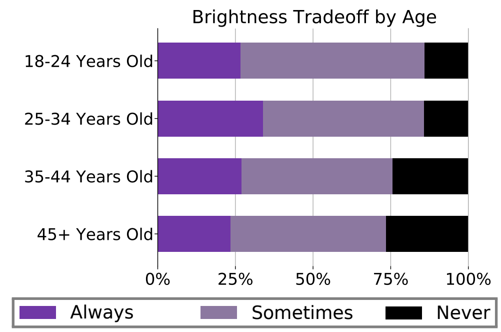

Figure 4: How likelihood to accept a *brightness tradeoff* varied by the participant's age.

as P219 saying they would not accept reduced brightness when "trying to show the game to some near-blind gen xer like my dad."

For the 90 participants who responded *always*, many cited battery savings in their reason. Others mentioned that they generally prefer lowered screen brightness or already lower it on their device (24.4%), or that reduced brightness feels better on their eyes (20%). As P194 put it: "I always reduce screen blindness and utilise night mode, I use electronic devices frequently and do not like the sensation of burning retinas after midnight." Additionally, 22.2% of these participants felt that the reduced screen brightness would not affect their ability to perform their assigned task. For example, P161 felt "the dull screen is still very easy to see and appreciate."

Like participants who indicated *always*, participants who chose *never* to accept reduced screen brightness were also concerned with their eye health. Of these 54 participants, 18.5% mentioned concerns about eye strain in their explanations. Participants who chose *never* also reported a general preference for clear displays and full brightness (24.1% and 27.8% of participants, respectively). 18.5% of these participants were also concerned about being able to see well enough to complete their assigned task. Finally, 20.4% of participants felt the reduced brightness tradeoff was not worth it, such as P30: "I prefer a bright, clear screen. 20 min wouldn't be enough to deal with the lowness of the other."

The 160 participants who selected *sometimes* also differed on whether reduced or normal screen brightness was better for their eyes. When asked for situations where participants would accept the tradeoff, nine mentioned eye protection and four preferred reduced brightness anyways. Conversely, when asked for situations where participants would not accept the tradeoff, nine said that the normal brightness was easier on their eyes, and eight preferred normal brightness most of the time.

Besides battery characteristics, the most common factor that affected acceptance *sometimes* was the light level of the surrounding environment. 17.5% of these participants would accept reduced brightness in any dim lighting, 16.9% would accept it at nighttime,

Table 7: Our parsimonious ordinal logistic regression model for respondents' willingness to accept a *GPS tradeoff*.

Factor	Baseline	OR	Coeff.	SE	z	p
Context: Gaming	Navigation	1.398	0.335	0.395	0.849	0.396
Context: Photo	Navigation	3.037	1.111	0.415	2.680	**0.007**
Quality Level	(linear fit)	2.194	0.786	0.293	2.685	**0.007**
Gender: Woman	Man	0.472	-0.752	0.347	-2.165	**0.030**
Gender: Non-binary	Man	1.005	0.005	1.033	0.005	0.996

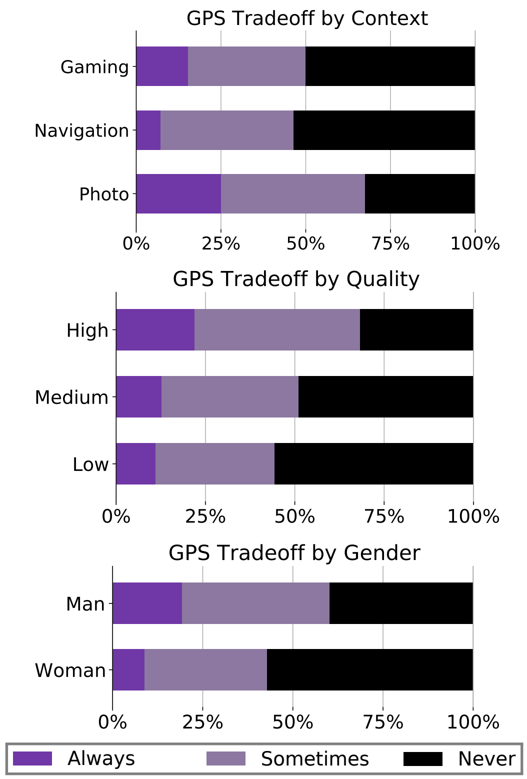

Figure 5: How participants' likelihood to accept a *GPS tradeoff* varied across the three significant dimensions. Because only a few participants identified as non-binary, the sample size is too small for meaningful comparison.

and 12.5% would accept it when indoors. Additionally, 28.8% mentioned not accepting reduced brightness when in bright light or sunlight, 22.5% when outdoors, and 8.1% during the daytime.

4.4.3 GPS. In our regression (Table 7), we found that participants' preferences about the GPS tradeoff varied by context, quality level, and the participant's gender. Compared to the navigation context, participants were 3.0× as likely to accept the GPS tradeoff in the photo context (p = .007), while preferences did not differ significantly in the gaming context. Note that GPS was not directly applicable in the other three contexts, so we did not ask those participants about GPS. With an increase in quality level and a corresponding decrease in battery savings, participants were 2.2× as likely to accept the GPS tradeoff (p = .007). Finally, compared to men, women were 0.47× as likely (i.e., less likely) to accept the GPS tradeoff (p = .030). Figure 5 visualizes these trends.

For the 54 participants who would *sometimes* accept the GPS tradeoff, the most common factors that affected acceptance were characteristics of the navigation route and battery status. Additionally, 18.5% of these participants mentioned that they would not accept reduced GPS accuracy when playing a game that uses location data, and 18.5% mentioned that they would not accept the reduction for any kind of navigation situation. While these participants selected that they would *sometimes* accept the tradeoff, their free-response answers indicate that they actually would never accept the tradeoff anytime it actually affected them.

This trend was common across all of the GPS responses. Of the 21 participants who selected that they would *always* accept reduced GPS accuracy for their assigned task, only 6 (28.6%) indicated that they would still always accept the tradeoff in other situations. Out of the 66 participants who said they would *never* accept the tradeoff for their assigned task, only 22 (33.3%) were willing to accept the tradeoff in other circumstances, and only 6 of these participants actually named contexts that were potentially relevant to GPS accuracy. The remaining situations were: when using apps that do not need exact location data (9 participants), when not using their phones to navigate (8), or when using their phone for social media or web browsing, both tasks that do not typically use frequently updated location data (6).

Participants who would *never* accept the GPS tradeoff most commonly expressed a need for accurate location information. Out of 66 participants, 40.9% expressed a sentiment similar to P71, who wrote, "I would want to know precise information about my location at all times." More specifically, 28.8% were concerned with getting lost or otherwise messing up their navigation, such as P301: "The reduced accuracy would cause me to get lost and miss exits due to the lag of the GPS system behind my movements." Finally, 18.2% felt that the reported increase in battery was not enough to make the tradeoff worth it, such as P73: "It barely gives any extra time to the battery, and it would be really annoying to deal with."

4.5 Relative Preferences About Tradeoffs

For three contexts, all three tradeoffs applied, so we first ran an omnibus Friedman test, subsequently running pairwise tests if the omnibus test was significant. Participants' willingness to accept the tradeoffs varied significantly in the game context (Friedman $\chi^2(2) = 11.757$, $p = .003$). In particular, willingness to accept reduced GPS was less than the willingness to accept reduced brightness or speed, with marginal significance (both p = .056). Participants' willingness to accept the tradeoffs also varied significantly in the navigation context (Friedman $\chi^2(2) = 39.181$, $p < .001$). In particular, participants were significantly less willing to accept reduced GPS than to accept reduced brightness or speed (both $p < .001$). We did not observe significant differences in the photo context (Friedman $\chi^2(2) = 1.219$, $p = .544$).

In the remaining three contexts, GPS did not have an obvious application to the app scenario, so we only asked about brightness and speed tradeoffs. Participants were more likely to accept reduced speed than reduced brightness in both the video (Friedman $\chi^2(1) = 5.121$, $p = .024$) and web browsing (Friedman $\chi^2(1) = 9.000$, $p = .003$) contexts. In both cases, roughly twice as many participants were *always* willing to accept the speed tradeoff than were

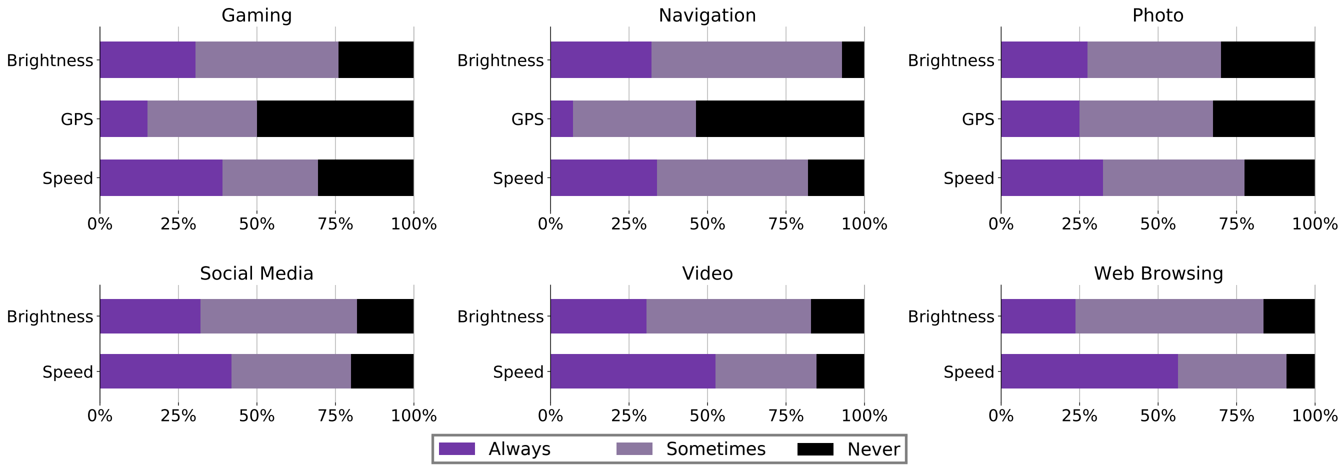

Figure 6: Each participant answered questions about brightness, GPS, and speed tradeoffs for their single assigned context and quality level. This graph shows how preferences varied across tradeoffs within each context. Note that participants in the social media, video, and web browsing context were not asked about GPS, which was not applicable to the app scenario used.

always willing to accept the brightness tradeoff. In contrast, we did not observe significant differences in the social media (Friedman $\chi^2(1) = 0.154$, $p = .695$) context. Figure 6 shows how preferences for tradeoffs varied within a context.

5 DISCUSSION

We conducted a 304-participant online survey with a mixed-subjects design to study contextual preferences for tradeoffs between power savings and performance. Among the six contexts studied, participants were most likely to accept speed and brightness reductions when browsing the web or using social media. Reasonably, participants were least likely to accept performance reductions that would heavily hinder the task: GPS for navigation, speed for gaming, and brightness for taking photos. Within each context, participants preferred reductions in brightness or speed over GPS degradation. Participants also preferred speed reductions over brightness reductions when watching videos or browsing the web.

Our results suggest that phones should consider usage context when managing power and performance, which existing smartphones do not offer. A context-aware and self-aware system for power-management would fill this gap while being minimally intrusive to the user. We make the following recommendations for implementing such a system.

Leverage Predictive Personas. We found six personas that capture a large set of user preferences. Self-aware systems can leverage these personas to efficiently predict user preferences about power and performance tradeoffs. We envision that the smartphone can ask a few short questions that allow it to assign the user to any relevant personas, subsequently adjusting power and performance settings accordingly. For example, the phone can ask the user whether they care about their phone performance while playing games in order to determine whether gaming is an appropriate context for automatically lowering the processing speed or brightness.

Investigate How to Predict Power Goals. Across all tradeoffs, participants were concerned about their smartphone's current battery level, when they will be able to charge their phone next, and their predicted battery needs until then. Future work should research accurate means for predicting the latter events, through a combination of user input and usage data. For example, the system could ask the user to predict whether their upcoming day will differ from their norm in terms of workload or charger access, and then determine power goals for the day using its understanding of a typical day's usage and charging patterns.

Utilize Sensor Information. Power considerations should not be isolated to the user and app context alone. For example, our participants indicated environmental light level was important in their decision regarding the brightness tradeoff. The smartphone should consider the user's ideal screen brightness for a given ambient light in addition to their personas. As many phones already adjust screen brightness based on user preferences and sensed illuminance in the environment [5, 34], this particular feature would be easy to incorporate into a context- and self-aware system.

Pay Careful Consideration to Navigation. Participants often expressed hesitation towards performance degradation during navigation, but also felt uncomfortable having low battery while navigating an unfamiliar place. This tension highlights the need for the self-aware system to specifically consider GPS navigation needs when setting power goals. Future work can investigate whether this can be done automatically or through asking the user directly.

ACKNOWLEDGMENTS

We thank Anne Farrell and Will Kong for their assistance in modeling the tradeoffs we studied.

REFERENCES

[1] Eyal Amir, Michael Anderson, and Vinay Chaudhi. 2007. *Report on DARPA Workshop on Self-Aware Computer Systems.* Technical Report. SRI International. https://apps.dtic.mil/sti/citations/AD1002393
[2] Android Developers. 2020. Power Management. https://developer.android.com/about/versions/pie/power
[3] Apple. 2018. Get the Top Apps of 2018. https://apps.apple.com/story/id1438574124
[4] Apple. 2019. Use Low Power Mode to Save Battery Life on Your iPhone. https://support.apple.com/en-us/HT205234

[5] Apple. 2020. Adjust the Brightness on Your iPhone, iPad, or iPod Touch. https://support.apple.com/en-us/HT202613
[6] Kumaripaba Athukorala, Eemil Lagerspetz, Maria von Kügelgen, Antti Jylhä, Adam J. Oliner, Sasu Tarkoma, and Giulio Jacucci. 2014. How Carat Affects User Behavior: Implications for Mobile Battery Awareness Applications. In *Proc. CHI*.
[7] Avast. 2017. Avast Android App Performance and Trend Report Q3 2017. https://press.avast.com/press-kits/android-trend-report-q3-2017
[8] AVG Technologies. 2016. Android App Performance and Trend Report H1 2016. https://cdn2.hubspot.net/hubfs/2683519/signal_pdfs/avg-technologies-android-app--performance-trends-report-h1-2016.pdf
[9] Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. 2007. A Survey on Context-Aware Systems. *IJAHUC* 2, 4 (June 2007), 263–277.
[10] Nilanjan Banerjee, Ahmad Rahmati, Mark D. Corner, Sami Rollins, and Lin Zhong. 2007. Users and Batteries: Interactions and Adaptive Energy Management in Mobile Systems. In *Proc. UbiComp*.
[11] Jose Felipe Mejia Bernal, Luca Ardito, Maurizio Morisio, and Paolo Falcarin. 2010. Towards an Efficient Context-Aware System: Problems and Suggestions to Reduce Energy Consumption in Mobile Devices. In *Proc. ICMB-GMR*.
[12] comScore. 2017. The 2017 U.S. Mobile App Report. https://www.comscore.com/Insights/Presentations-and-Whitepapers/2017/The-2017-US-Mobile-App-Report
[13] Anind Dey, Gregory Abowd, and Daniel Salber. 2001. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. *Human-Computer Interaction* 16 (2001).
[14] Amandeep Dhir, Puneet Kaur, Nobert Jere, and Ibrahim A. Albidewi. 2012. Understanding Mobile Phone Battery - Human Interaction for Developing World - A Perspective of Feature Phone Users in Africa. In *Proc. BCFIC*.
[15] Rob Eisinga, Tom Heskes, Ben Pelzer, and Manfred Te Grotenhuis. 2017. Exact P-values for Pairwise Comparison of Friedman Rank Sums, With Application to Comparing Classifiers. *BMC Bioinformatics* 18, 1 (2017).
[16] Denzil Ferreira, Anind K. Dey, and Vassilis Kostakos. 2011. Understanding Human-Smartphone Concerns: A Study of Battery Life. In *Proc. Pervasive*.
[17] Denzil Ferreira, Eija Ferreira, Jorge Goncalves, Vassilis Kostakos, and Anind K. Dey. 2013. Revisiting Human-Battery Interaction with an Interactive Battery Interface. In *Proc. UbiComp*.
[18] Dan Frakes, Nick Guy, and Kevin Purdy. 2016. What You Should (and Shouldn't) Do to Extend Your Phone's Battery Life. New York Times. https://www.nytimes.com/wirecutter/blog/what-you-should-and-shouldnt-do-to-extend-your-phones-battery-life/#battery-saving-myths
[19] David Garlan, Daniel P. Siewiorek, Asim Smailagic, and Peter Steenkiste. 2002. Project Aura: Toward Distraction-Free Pervasive Computing. *IEEE Pervasive Computing* 1, 2 (2002), 22–31.
[20] Benjamin Gaudette, Carole-Jean Wu, and Sarma Vrudhula. 2016. Improving Smartphone User Experience by Balancing Performance and Energy With Probabilistic QoS Guarantee. In *Proc. HPCA*.
[21] Google. 2020. Get the Most Life From Your Android Device's Battery. https://support.google.com/android/answer/7664692?hl=en
[22] Todd Haselton. 2017. Stop Quitting Apps on Your iPhone and iPad…It's Making Things Worse. CNBC. https://www.cnbc.com/2017/07/20/stop-quitting-apps-on-your-iphone-and-ipad-its-making-things-worse.html
[23] Chi-Kai Ho, Chung-Ta King, and Yung-Ju Chang. 2019. SCQ: Stage-Based, Context-Aware, QoE-Driven Power Optimization for Interactive Applications on Mobile Devices. In *Proc. MDM*.
[24] Henry Hoffmann. 2015. JouleGuard: Energy Guarantees for Approximate Applications. In *Proc. SOSP*.
[25] Henry Hoffmann, Jim Holt, George Kurian, Eric Lau, Martina Maggio, Jason E. Miller, Sabrina M. Neuman, Mahmut Sinangil, Yildiz Sinangil, Anant Agarwal, Anantha P. Chandrakasan, and Srinivas Devadas. 2012. Self-Aware Computing in the Angstrom Processor. In *Proc. DAC*.
[26] Henry Hoffmann, Martina Maggio, Marco Santambrogio, Alberto Leva, and Anant Agarwal. 2010. *SEEC: A Framework for Self-aware Computing*. Technical Report. MIT CSAIL.
[27] Simo Hosio, Denzil Ferreira, Jorge Goncalves, Niels van Berkel, Chu Luo, Muzamil Ahmed, Huber Flores, and Vassilis Kostakos. 2016. Monetary Assessment of Battery Life on Smartphones. In *Proc. CHI*.
[28] Andreas Krause, Asim Smailagic, and Daniel P. Siewiorek. 2006. Context-Aware Mobile Computing: Learning Context-Dependent Personal Preferences From a Wearable Sensor Array. *IEEE Transactions on Mobile Computing* 5, 2 (2006), 113–127.
[29] Wooseok Lee, Reena Panda, Dam Sunwoo, Jose Joao, Andreas Gerstlauer, and Lizy John. 2018. BUQS: Battery- and User-Aware QoS Scaling for Interactive Mobile Devices. In *Proc. ASP-DAC*.
[30] Peter Lewis, Arjun Chandra, Funmilade Faniyi, Kyrre Glette, Tao Chen, Rami Bahsoon, and Jim Torresen. 2015. Architectural Aspects of Self-Aware and Self-Expressive Computing Systems: From Psychology to Engineering. *Computer* 48 (2015), 62–70.
[31] LG Electronics USA. 2016. "Low Battery Anxiety" Grips 9 Out Of Ten People. *PR Newswire* (May 2016). https://www.prnewswire.com/news-releases/low-battery-anxiety-grips-9-out-of-ten-people-300271604.html
[32] Sakkthivel Manickam, Moovendhan Vellasamy, and Githa Heggde. 2020. Investigating the Relationship Between Age and Smart Phone Usage Patterns: Evidences from Indian Smart Phone Users. *International Journal of Business Excellence* 21 (2020), 394–409.
[33] Marcelo Martins and Rodrigo Fonseca. 2013. Application Modes: A Narrow Interface for End-User Power Management in Mobile Devices. In *Proc. HotMobile*.
[34] Emil Protalinski. 2018. How Android Pie's Adaptive Battery and Adaptive Brightness Work. VentureBeat. https://venturebeat.com/2018/08/28/how-android-pies-adaptive-battery-and-adaptive-brightness-work/
[35] Ahmad Rahmati, Angela Qian, and Lin Zhong. 2007. Understanding Human-Battery Interaction on Mobile Phones. In *Proc. MobileHCI*.
[36] Nishkam Ravi, James Scott, Lu Han, and Liviu Iftode. 2008. Context-Aware Battery Management for Mobile Phones. In *Proc. PerCom*.
[37] Bill Schilit, Norman Adams, and Roy Want. 1995. Context-Aware Computing Applications. In *Proc. WMCSA*.
[38] Hao Shen, Qiuwen Chen, and Qinru Qiu. 2014. Battery Aware Stochastic QoS Boosting in Mobile Computing Devices. In *Proc. DATE*.
[39] Khai N. Truong, Julie A. Kientz, Timothy Sohn, Alyssa Rosenzweig, Amanda Fonville, and Tim Smith. 2010. The Design and Evaluation of a Task-Centered Battery Interface. In *Proc. UbiComp*.
[40] Jian Xu, Syed Masum Billah, Roy Shilkrot, and Aruna Balasubramanian. 2019. DarkReader: Bridging the Gap Between Perception and Reality of Power Consumption in Smartphones for Blind Users. In *Proc. ASSETS*.
[41] Dongsong Zhang, Boonlit Adipat, and Yaser Mowafi. 2009. User-Centered Context-Aware Mobile Applications – The Next Generation of Personal Mobile Computing. *Communications of the Association for Information Systems* 24 (2009).