Good Vibrations: Can a Digital Nudge Reduce Digital Overload?

Fabian Okeke[†], Michael Sobolev[†Φ], Nicola Dell[†Ψ], Deborah Estrin[†]
[†]Cornell Tech, [Φ]Technion, [Ψ]The Jacobs Institute
[fno2;ms3377;nixdell;de226]@cornell.edu

ABSTRACT

Digital overuse on mobile devices is a growing problem in everyday life. This paper describes a generalizable mobile intervention that combines nudge theory and negative reinforcement to create a subtle, repeating phone vibration that nudges a user to reduce their digital consumption. For example, if a user has a daily Facebook limit of 30 minutes but opens Facebook past this limit, the user's phone will issue gentle vibrations every five seconds, but the vibration stops once the user navigates away from Facebook. We evaluated the intervention through a three-week controlled experiment with 50 participants on Amazon's Mechanical Turk platform with findings that show daily digital consumption was successfully reduced by over 20%. Although the reduction did not persist after the intervention was removed, insights from qualitative feedback suggest that the intervention made participants more aware of their app usage habits; and we discuss design implications of episodically applying our intervention in specific everyday contexts such as education, sleep, and work. Taken together, our findings advance the HCI community's understanding of how to curb digital overload.

CCS CONCEPTS

•**Human-centered computing** → *Ubiquitous and mobile computing design and evaluation methods;*

KEYWORDS

Digital Overload; Social Media; Vibration; Intervention; Digital Nudge; Negative Reinforcement; Smartphones

ACM Reference format:
Fabian Okeke[†], Michael Sobolev[†Φ], Nicola Dell[†Ψ], Deborah Estrin[†]. 2018. Good Vibrations: Can a Digital Nudge Reduce Digital Overload?. In *Proceedings of MobileHCI '18: 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, Barcelona, Spain, September 3–6, 2018 (MobileHCI '18),* 11 pages.
DOI: 10.1145/3229434.3229463

1 INTRODUCTION

The amount of time that people spend consuming digital content has dramatically increased over the years because of the opportunity to use mobile devices in nearly every context and at nearly every moment. Popular mobile apps such as Facebook, YouTube, Snapchat, and more, have become an integral part of people's everyday lives, helping them to share their thoughts, connect with friends and families, receive news updates, and enjoy many forms of digital entertainment. However, research has shown that technology burdens people with the pressure of continual availability [2, 37], the need to constantly check in [38, 47], and the ease to procrastinate on work, studies, and personal goals [22, 39, 52]. As a result, many people desire to reduce or limit their technology use [21]. A number of studies have suggested completely abandoning social media [3], using feature phones instead of smartphones [31], calling users daily to reflect on their social media habits [6], limiting usage through personal productivity tools [16, 25], and committing as a group to reduce digital distraction [29]. Many of these existing solutions combine a variety of different interventions at the same time, making if difficult to understand the effects of particular behavioral theories or techniques.

The goal of our research is to design, implement, and evaluate a generalizable technique that reduces digital consumption on mobile devices. Specifically, we created a mobile app intervention that uses nudge theory [53] from behavioral economics and negative reinforcement [20, 51] from behavioral psychology. A nudge refers to an intervention that steers people in a particular direction without eliminating their freedom of making the final choice; while negative reinforcement refers to the strengthening of a behavior by avoiding a negative outcome or aversive stimulus. We combined these two concepts to design our digital intervention as a subtle, repeating phone vibration that nudges a user to stop using a target mobile application whenever they exceed a daily usage limit.

We evaluated our intervention through a controlled experiment with 50 workers recruited on Amazon's Mechanical Turk (MTurk) platform. Participants were randomly assigned to one of three experimental conditions: *control*, *rigid* and *personalized*, and each group used the intervention for three weeks. Although we designed our intervention to work on any target mobile application, we conducted our evaluation only on participants' Facebook usage in order to avoid experiment confounds. In the *control* condition, no intervention was applied; in the *rigid* condition, the daily Facebook usage limit was pre-determined and remained static over the study period; while in the *personalized* condition, a daily Facebook usage limit was calculated based on each participant's usage habits. During the first week of the experiment we did not deliver any explicit interventions in order to measure each participant's baseline usage. In the second week, participants in the *rigid* and *personalized* conditions experienced the intervention whenever they used Facebook past their daily limit. In the final week, we monitored participant's usage after withdrawing the intervention. We hypothesized that participants who had experienced the intervention conditions would use Facebook less than participants in the *control* condition. We also hypothesized that participants in *personalized*

MobileHCI '18, Barcelona, Spain

DOI: 10.1145/3229434.3229463

condition would show an increased reduction in usage than those in the *rigid* condition.

Our findings show that combining nudges and negative reinforcement successfully reduced participants' digital overload. We found statistically significant differences between the three experimental conditions. Participants in the *rigid* and *personalized* conditions changed their behavior but not participants in the *control* condition. Our results show that providing only textual feedback to participants about their daily digital consumption is not sufficient to change behavior. However, nudging participants via subtle, repetitive vibrations reduced their digital consumption by over 20% daily during the intervention week. When the intervention was removed, participants returned to their baseline behavior at the beginning of the experiment. Although, we hypothesized that the *personalized* condition would have a greater effect than the *rigid* condition, we did not find any detectable differences between these strategies. These findings suggest that subtle, repetitive vibrations are an effective way to reduce digital overload on mobile phones.

Our intervention offers several potential benefits for researchers and practitioners. It can be applied contextually: to nudge people back to their social interactions when they become distracted using their devices, to persuade individuals to stop using their phones in bed when it is time to sleep, to encourage individuals to return back to work when they become distracted by their mobile devices, and to nudge students in educational settings to use social media less during exam weeks when intense focus is necessary. The intervention is easy to adopt because it is a low-cost technique that does not require any additional equipment—vibration capabilities are ubiquitously available in modern-day digital devices.

2 BACKGROUND AND RELATED WORK

There has been growing concern within the HCI community about the negative effects of technology overuse and different solutions have been studied. Baumer et al. [3] investigated how people completely abandon social media and the challenges they face after making this decision. Researchers have encouraged the use of feature phones in place of smartphones [31], placed daily phone calls that asked users to reflect on their social media habits [6], and designed a standalone app for users to commit as a group to collectively reduce mobile phone distractions in social gatherings[29]. Hiniker et al. [16] applied daily goal setting, pop-up messages, and realtime feedback to reduce daily mobile phone usage. However, these approaches combine multiple interventions to reduce digital overload making it difficult to understand the effects of particular behavioral theories or techniques.

Researchers have shown that self-awareness can lead to changes in people's performance or behavior [5]. To increase self awareness, HCI researchers have designed several self-monitoring systems. Ubifit [8] uses a glanceable display in the background of mobile phones to increase users' self-awareness of their physical activities and persuade them to become more active. MyTime [16] provides a status notification bar that shows a user the amount of time spent on a current mobile app and the total time spent on multiple monitored apps. Kim et al. [25] framed time spent on desktop devices as positive or negative feedback to make users conscious of their digital habits. Several productivity tools (such as RescueTime, Focus, Moment, UnGlue) provide daily feedback that summarize users' digital habits so users can regulate their digital consumption. In addition to providing feedback, HCI researchers have studied the regulation of overuse through application blocking on an individual level [24, 33], in groups [28, 29], as a family [27], and in a classroom [23]. Our work extends the literature by investigating an approach that does not block users' access to services but nudges them to stop technology overuse.

Our study combines nudges [53] from behavioral economics and negative reinforcement [20, 51] from behavioral psychology to design a digital intervention that persuades users to spend less time using a target mobile app. The first concept, nudges, refer to interventions that steer people in a particular direction without eliminating their freedom of making the final choice. In this paper, we use *digital nudge* described by Okeke et al. [42] as *"nudges that are provided via digital technologies. Digital nudges can provide information, reminders and planning prompts to the users in the form of status-bar messages, pop-ups, phone vibration, and phone LED display."* We chose a continual gentle vibration as our digital nudge because vibration is considered private and subtle [15]. HCI researchers have studied the efficacy of vibration in supporting stroke survivors [40], to inform users of their proximity to places of interest during the exploration of cities [18], and to help users gain perspective of the objects of an interactive map [46]. Our work extends existing research by applying vibration to curb digital overload.

The second concept, negative reinforcement, refers to the strengthening of a behavior by avoiding a negative outcome or aversive stimulus. This process involves behavioral learning based on personal experiences over time and it provides a high potential for successful behavior change [12, 13]. One common example of negative reinforcement is when a driver starts a car without putting on the seat belt. This leads to a repeating beep sound in the car until the seat belt is worn to stop the irritating sound. When the driver enters the car in the future, the seat belt is immediately worn due to learned behavior to avoid the aversive beeping sound. Another example is when an individual's phone vibrates every time a distracting app is used but vibration ceases immediately the user stops using the distracting app. Application of negative reinforcement has successfully reduced cheating in college exams and increased enforcement of safety rules in factories and hospitals [12]. However, there is little work on applying this concept to reduce digital overload.

3 INTERVENTION DESIGN

The goal of our research is to design a mobile intervention that reduces digital overload on mobile devices. To achieve this, we designed a standalone Android[1] application with two design features: (1) realtime, textual feedback that tells a user the amount of time they have spent on a mobile application, and (2) a digital nudge, using gentle phone vibration, when usage of a mobile app surpasses a daily time limit.

We adopted a subset of Consolovo's design principles [9] for behavior change technologies in designing our realtime feedback.

[1]We focused on Android only because iOS does not provide programmatic access to monitor user applications.

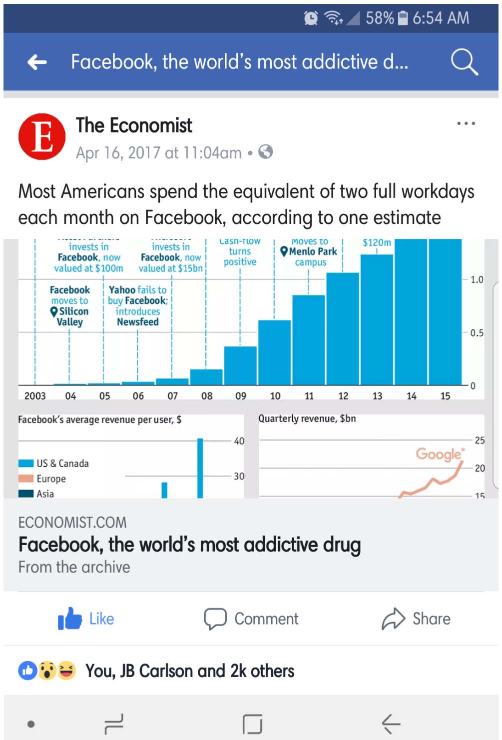

Figure 1: In this scenario, the user engages with Facebook without checking how much time has been spent or how many times it has been opened in the current day. The user can engage with Facebook without any interruption from the status bar feedback that is currently hidden because it is designed to be ignorable.

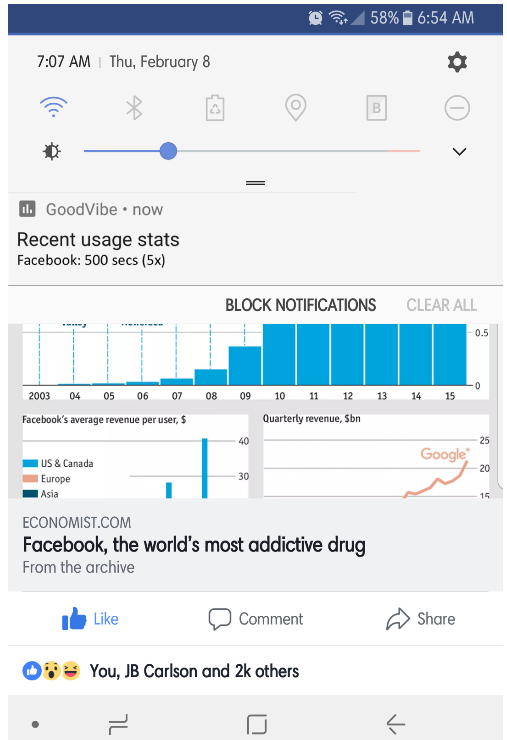

Figure 2: In this scenario, the user swipes down on the status bar to see that 500 seconds has been spent on Facebook and it has been opened five times in the current day. All users can see the status bar but only users in intervention group experience gentle vibrations when Facebook daily limit is surpassed.

These principles explain that a persuasive technology should be *abstract and reflective*, *unobtrusive*, and *support occasional ignorability*. To achieve an *abstract and reflective* design, a persuasive system should show summarized feedback to users instead of raw data collected and show users how the provided information relate to their behavior; it should *support occasional ignorability* by enabling users to ignore the technology when they chose to; and it should be *unobtrusive* by collecting data without unnecessarily interrupting the user.

We designed our intervention app to be *abstract and reflective* by providing a status bar feedback that summarizes how long a target application has been used and how many times it had been opened in a day. Figures 1 and 2 show a Facebook view before and after looking at the status bar information. This information is updated every five seconds thereby providing continuous feedback. To make the app *ignorable*, we designed the persistent status bar notification to be low-priority on android so that users can ignore the ongoing feedback or swipe down the status bar to access it. As a low-priority notification, it protects a user's privacy from prying eyes because the notification is not visible when the phone screen is locked. To achieve the last design principle, all data was collected in an *unobtrusive* manner without interrupting a user's daily phone operations.

To create the second design feature of administering a digital nudge, we designed the intervention app to nudge a user with subtle phone vibrations whenever the target mobile app was used past a specified limit. We designed the subtle vibration to be negatively reinforcing by repeating it every five seconds similar to Pielot et al. [44] who tested multiple vibration pulses and found that a five-second pause prevents quick adaptation[2]. For example, if a user has a daily Facebook limit of 30 minutes but opens Facebook past this limit, the user's phone will issue gentle vibrations every five seconds, but the vibration stops once the user navigates away from Facebook. We chose vibration as our digital nudge because it is considered private and subtle [15] and it is ubiquitously available on all mobile devices. HCI researchers have also found evidence that users can consume a gentle continual vibration pulse without becoming annoyed or distracted and they can detect within a reasonable amount of time when the vibration stops [44]. Therefore, vibration has potential as a suitable nudge to remind users of their digital habits.

After designing and implementing the intervention app, we pilot tested it with 10 participants recruited from Amazon's Mechanical

[2]In designing the vibration feedback, we wanted it to be gentle yet noticeable so we programmed the vibration cue to have two pulses that are each 100 ms long, and 100 ms apart (buzz first for 100 ms, wait for 100 ms, and finally buzz for 100 ms).

Turk (MTurk) platform. Participants installed and used the app for five days, which helped us address technical bugs and confirm that the system was working as expected before deploying the full experiment described in the next section. The deployed version of our intervention app is freely available on Github as open source software[3].

4 EXPERIMENT DESIGN

To evaluate the impact of digital nudges and negative reinforcement on digital overload, we conducted a controlled experiment with 50 workers recruited through Amazon's Mechanical Turk (MTurk) platform—a crowdsourcing platform where workers complete tasks such as surveys and categorization in exchange for small monetary payments[1]. A growing number of HCI studies recruit participants from MTurk [17, 26, 30] because researchers can access a large pool of geographically distributed populations at a relatively low cost.

Our experiment examined how subtle phone vibrations reduced the amount of time participants spent using a mobile application. We carefully setup the experiment to avoid study confounds. Although we implemented our intervention app to monitor usage of any mobile application, we selected only a single target application for all participants in order to avoid confounds that may result from allowing each participant to choose a wide variety of different—and potentially unknown, unreliable, or buggy—apps from the PlayStore. We chose Facebook as the target application because it is the most dominant mobile application in the world [11] and prior research has shown that a large proportion (e.g. 90%) of people who are interested in limiting their consumption of digital content indicate that they would like to reduce the amount of time spent using Facebook [16].

Our experiment ran for three weeks: a one-week baseline period, a one-week intervention period, and a one-week follow-up period. We assigned a limit to each participant both for the amount of time spent on Facebook and the number of times Facebook was opened in a day. We did not allow participants to choose their own limits because different levels of engagement could confound the experiment. For example, some participants may be more motivated to set their daily limits than others. All participants were randomly assigned into one of three conditions that determined their daily Facebook limits: *control*, *rigid*, and *personalized*.

Control: Participants in this condition did not receive any limits on daily time spent and number of opens of Facebook. As such, they did not experience any vibration feedback throughout the duration of the experiment.

Rigid: Participants in this condition received a limit that was informed by how people felt daily Facebook use should be regulated. We asked 100 participants on MTurk the maximum time that should be spent on Facebook and the maximum number of times it should be opened per day on mobile phones. Median response was 30 minutes and five times per day with one-third of the users indicating at most two opens per day (mean: 49.4 minutes, 11 opens; this implies 4.5 minutes per open). Based on these responses, we set the limit as two opens per day each five minutes long so that the majority of participants in this condition would experience vibration regardless of their personal usage habits.

Personalized: Participants in this condition received limits that were based on their personal usage. We developed the app as a small data application [19]—an application that derives behavioral insight from a user's past habits—by computing each participant's daily limit based on their behavior in the first week. Then we set users' daily limits to half (50%) of the average time they spent in the baseline week. For example, suppose on average, a participant used Facebook for one hour daily and opened it ten times per day throughout the baseline week then the participant's limit during the intervention week would be set to half of their usage: 30 minutes and five times per day. Since this is the first study exploring the application of subtle vibrations to reduce digital overload, we decided to begin by nudging users towards spending half the time they spend becoming digitally overloaded and potentially adjusting this ration in future work[4].

Participants in all three experimental conditions had the status bar notification that showed realtime feedback of app usage. We kept this consistent across all participants for two reasons: (1) we wanted differences in vibration strategies to be the only difference across conditions; and (2) the experiment on-boarding process explained to participants that the study was about understanding their Facebook usage habits. By making all users aware of the purpose of the study, it was easier to isolate and test the effect of vibration on specific groups of participants. Moreover, popular apps (e.g. RescueTime, Moment, Focus) for combating digital overload provide users with feedback about their digital habits. As such, using the usual behavior as baseline better reflects current practices and improves the ecological validity of our findings.

Study Procedure

A key difference between our study and prior work on digital overload [16, 25, 29] is that we recruited all participants including individuals who did not self-identify as people willing to reduce their Facebook usage. Since we focused our recruitment on MTurk, we initially considered asking MTurk workers if they were interested in reducing their Facebook usage and subsequently filtering only qualified users. However, we avoided this approach because it would lead to demand characteristics [43], a situation where participants anticipate the purpose of our study and aim to be "good" users that would not "ruin" the study—they would claim to be interested and may intentionally avoid using Facebook during the intervention week leading to false effects. As such, we advertised the study as research that investigated participants' Facebook usage habits and nudged them through vibrations whenever they used Facebook past a daily assigned limit.

To participate in the study, MTurk workers had to be active Facebook users (at least five minutes of usage per day). We filtered eligible users from a general survey about technology habits that asked MTurk workers to estimate how much time they spent on Facebook on their mobile phones. Participants who reported using at least five minutes daily were contacted and asked to enroll in the study. During recruitment, we noticed that majority of our prospective participants (approx. 64%) were located in India while 31% were located in the United States and this could lead to study confounds. For example, festive holidays in India could lead to

[3] Vibration app: https://github.com/fnokeke/FgApp

[4] We initially designed two types of personalized strategies: one based on a single individual's daily usage time and the other based on the average of a group's daily app usage time. However, both yielded similar results.

different Facebook usage patterns among Indian users compared to United States users. To minimize this problem, we decided to conduct our experiment with only recruited users based in India because we had more participants in this location and globally, India has the highest number of Facebook users [32].

All enrolled participants were required to have Facebook mobile app already installed on their phones. We added this constraint because our intervention app can only monitor when users use mobile apps in the foreground but it cannot monitor when a service is accessed through a mobile browser (e.g. visiting facebook.com) because it is technically infeasible to implement this functionality on android devices. For this reason, we only recruited participants who already had Facebook mobile application installed on their phones prior to the study. After invited MTurk workers consented to the study, they followed instructions to downloaded the intervention app. Upon successful download, the app automatically checked if Facebook mobile app was already installed on participants' phones. If this was not the case, participants received an error stating that their phones were not compatible with the study and they could not continue. We did not ask participants to install Facebook mobile app because anyone who previously did not have it may likely not engage with it after installation.

After being randomized into one of the three experimental conditions, each participant took part in the study for a total of three weeks. The first week of the study consisted of a *baseline* period, in which we gathered data regarding participants application usage habits and provided feedback via usage statistics, but no vibration. The second week of the study was the *treatment* period, in which participants received both usage statistics and vibration feedback according to the experimental condition to which they were assigned. The third week of the study consisted of a *follow-up* period, in which any vibration feedback that the participant was receiving was removed, although the system continued to monitor the participant's daily usage and provide usage statistics. The goal of the follow-up week was to assess the effect on participants' behavior of removing the intervention. At the end of three weeks, all participants completed an exit survey and were instructed to uninstall the experimental system from their devices. Each participant received $15 (approx. 1000 rupees) for participating in the study and a lottery chance for an additional $5 bonus. For perspective, $1.00 can buy a 12-ounce cup of coffee in India and $5.00 is sufficient for a full meal.

Participant Demographics

We enrolled a total of 50 participants based in India (35 males and 15 females) via MTurk. Participants ranged in age from 19 to 40 years (average = 28.8, median = 29) and self-reported that they had worked on MTurk between one month and six years. Participants reported a diverse range of Facebook usage habits, self-reporting that they opened the Facebook application between one and 30 times per day, and spent between five minutes and three hours per day using Facebook. Participants used Facebook for several reasons including to see their friends' status updates, share their thoughts, interact with comments and pictures, and read the news.

Data Collection and Analysis

Our intervention app logged daily time spent on Facebook and the number of times it was opened in a day. In total, our system logged 42,389 server records for all participants across the entire three-week study. At the start of the study, we enrolled a total of 78 participants in the three experimental conditions: Control (n=25), Rigid (n=26), and Personalized (n=27). However, several participants did not complete the study so we restricted our analysis to only users who completed the experiment. This led to a total of 50 participants: *Control* (n=16), *Rigid* (n=19), and *Personalized* (n=15). We had two types of non-completers that we removed before performing our analysis: those without data for baseline week and those who had data for only the baseline week. The breakdown is as follows: control - (no baseline: 5, only baseline: 4); rigid - (no baseline: 5, only baseline: 2); personalized - (no baseline: 8, only baseline: 4). There were no indications that these users are different from those who completed the experiment as all participants went through the same recruitment process.

We analyzed the impact of our intervention on the total amount of time that each participant spent on Facebook for each day of the study and the total number of times Facebook was opened daily (see Table 1)[5]. Our analyses capture the change in application usage for each participant between a one-week *baseline* period, where participants across all experimental conditions did not receive gentle vibrations, and a one-week *treatment* period in which participants received the intervention according to the group to which they were assigned. Further, we capture the change for each participant during the one-week *follow-up* period, in which vibration feedback was removed for all experimental conditions.

Before performing our analyses, we evaluated the Kolmogorov-Smirnov test to assess the normality of our data and found that it was statistically significant for all of our outcome metrics. As a result, we used non-parametric Wilcoxon Signed-Rank tests to analyze changes within experimental conditions across study periods.

5 EXPERIMENT RESULTS

Our experimental results show that combining digital nudge with negative reinforcement to create a subtle, repetitive vibration, successfully reduced the amount of time that participants spent using the target application regardless of the strategy used to compute the daily app limit. However, the reduction in usage did not persist after the intervention was removed in the follow-up period. Findings from qualitative feedback that we gathered from participants shows that they found the realtime feedback useful for keeping track of their own usage, and that they correctly perceived the vibration as providing a reminder that encouraged them to spend less time using the target application. The rest of this section describes these findings in detail.

Vibration successfully reduces digital consumption of the target application.

Our findings show that our intervention led to a statistically significantly reduction in the amount of time participants spent using the target application—in this case, Facebook. Across all participants in the experiment, the maximum time spent on Facebook was 4.7 hours per day and the maximum number of times Facebook was opened daily was 99 times. Table 1 summarizes the results of Wilcoxon Signed-Rank tests and Cohen's d effect size that we conducted to analyze, for each experimental condition, the differences

[5]We do not include, in the table, details of statistical testing for the number of times Facebook was opened because they were non-significant.

Experiment Condition	Effect Size Treatment	p-value Treatment	Effect Size Followup	p-value Followup
Control	-0.081	0.510	0.010	0.316
Rigid	0.497	0.009**	0.293	0.138
Personalized	0.600	0.011*	0.137	0.262
All treatment	0.540	0.00***	0.227	0.108

Table 1: Average time spent on Facebook per day. Summary of pairwise comparisons between baseline and study periods show that for baseline-treatment comparison, there are statistically significant differences between the *rigid* and *personalized* conditions but not *control* condition.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

between the average amount of time that participants spent using Facebook across the *baseline* and *treatment* periods. Participants in *both* of the vibration conditions spent statistically significantly less time using the target application during the *treatment* period than in the *baseline* period. Figure 3 shows the average application usage for participants in each experimental condition during the three study periods while Figure 4 shows the percentage change in time spent across each group. The amount of time spent on Facebook per day during the *treatment* period decreased by over 20% for each of the vibration conditions: *Rigid* (22.8% or 8.5 minutes) and *Personalized* (23.3% or 8.5 minutes), but not for the *Control* condition, which remained at relatively same level throughout the experiment. These findings suggest that combining digital nudge with negative reinforcement is effective in reducing digital overload.

Participants perceived vibration as negative feedback. At the end of the experiment, participants were invited to complete an exit survey that provided qualitative information regarding their experiences during the study and opinions of the vibration application. We received a total of 31 survey responses: *Control* (n=10), *Rigid* (n=8), and *Personalized* (n=13). Several participants (19/31) who experienced the vibration pulses reported that it increased the awareness of their application usage. When asked how they responded when vibration started, 21 participants said that they stopped using the application after a few minutes, while the remaining 10 stated that they continued to use the application as normal. This finding supports our quantitative data, which shows decreased usage for participants in the vibration conditions. It also suggests that nudging users via gentle vibration does not always work. In addition to understanding how participant behavior changed when vibration started, we were also interested in understanding how the vibration made them feel. The majority (26/31) said that the vibration was irritating, with roughly half (n=14) saying they found it mildly irritating, and the other half (n=12) saying that they found it very irritating. This suggests that participants perceived the aversive nature of the negatively reinforced stimulus. Although the majority (17/26) of participants that found the vibration irritating said that it caused them to spend less time using the application, none of the participants said that they stopped using it completely. This indicates that participants' aversion to the vibration was not strong enough to cause them to leave the platform completely.

Realtime feedback increased awareness of digital habits. Of the 31 participants who completed the exit survey, 28 said that the usage statistics provided by the persistent notification bar increased their perceived self-awareness of how much time they spent using Facebook every day. Participants reported swiping down the status bar to look at their usage habits on average 12 times per day (median=9, mode=5). This result validates findings from a recent study by Rooksby et al. [50] that showed how personally tracking screen time on digital devices helped to improve overall awareness. Many participants submitted written comments suggesting that they enjoyed having the ability to keep track of their daily usage, with several describing how the feedback made them more aware of how much they were using Facebook: *"Through the status bar, I realized that I was using Facebook much more than what I had thought."*(P7). Several participants explained that receiving the feedback *"saved them time"* by causing them to use Facebook less than they had before. Taken together, these findings suggest that, regardless of whether the participants experienced gentle vibrations, the realtime feedback provided through the status bar was useful for participants and it increased their perceived awareness of personal digital habits.

Participants gradually return to their old habits when the intervention is removed
The goal of the *follow-up* period of the study was to investigate if any behavior change that occurred as a result of the vibration feedback persisted when the vibration was removed. We hypothesized that, in the absence of gentle vibrations, participants would gradually return to their original levels of application usage. As shown in Figure 4, the amount of time spent by participants in *rigid* and *personalized* increased between the *treatment* period, when they were experiencing vibration, and the *follow-up* period, when vibration was removed.

We conducted one-sided Wilcoxon Signed-Rank tests to analyze, for each experiment condition, the differences between the time spent on Facebook in the *baseline* period and in the *follow-up* period. The results of the tests were non-significant for all experimental conditions, indicating that there was no detectable difference between participants' usage of the application in the *baseline* and *follow-up* periods. This finding suggests that although receiving vibration effectively reduced application usage, there is no conclusive evidence that the effect may or may not be sustained after the intervention is removed.

There is no conclusive difference between rigid and personalized strategies.
In addition to understanding the overall impact of vibration, we were also interested in understanding if there were differences between the strategies used to compute the daily usage limits that would result in vibration. We hypothesized that personalizing the vibration to an individual participant's usage habits would be more effective and result in greater reductions in application usage. This hypothesis is based on prior work that suggests that personalization leads to more successful interventions by catering to each individual's context [14]. Although our analysis shows a descriptive difference in effect sizes (see Table 1) with the *personalized* condition having a higher effect size than the *rigid* condition, there is no statistically significant difference between both strategies. This result could be due to a number of factors. For example, in the

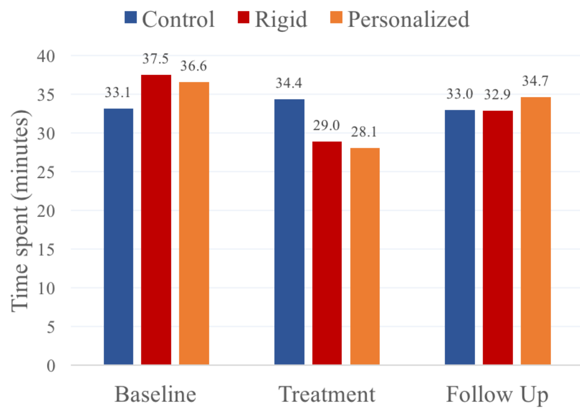

Figure 3: Time spent on Facebook across the three study periods. There is no difference between groups in the baseline week, however, there is a significant reduction in Facebook usage for the *rigid* and the *personalized* conditions during the treatment week but not for the *control* condition.

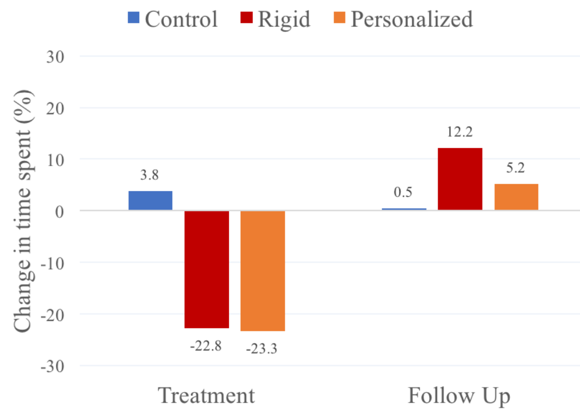

Figure 4: Change in time spent on Facebook relative to baseline week. During treatment week, daily average time spent significantly reduced by over 20% for the *rigid* and the *personalized* conditions but not for the *control* condition. However, the effect diminishes in the follow-up week.

personalized condition, we set the usage limit to be half (50%) of the participant's daily usage, which may not be the correct ratio to use. We also calculate the personalized usage limit based on participants' usage habit during the baseline week, but this period may not be long enough for determining a participant's true habit. Finally, we may not have enough statistical power to detect a significant difference between the groups given the small sample size of each group. Nevertheless, these limitations suggest exciting opportunities for more explorations. Future studies with a larger sample size and a longer study period will have more statistical power to detect and provide conclusive evidence on the comparison between personalized and static usage limits, which could eventually lead to the discovery of appropriate strategies for personalizing digital interventions to reduce digital overload.

Vibration does not reduce the number of times participants opened the target application.
In addition to measuring the total amount of time that participants used the target application in each experimental condition, we were interested in how the intervention may affect the number of times participants opened the target application. The results of Wilcoxon Signed-Rank tests that we conducted to analyze, for each experimental condition, the differences between the average number of times participants opened Facebook in the *baseline* period and the *treatment* period were non-significant for all three experimental conditions, indicating that the vibration did not detectably reduce the number of times that participants opened the application. This finding could be explained by the fact that participants only feel the effect of the intervention after they have already opened the Facebook app. As a result, even though a participant may open the application and immediately close it when they feel the vibration, their behavior will still 'count' as opening the application. This suggests that applying digital nudges to curb technology overuse may not be sufficient in regulating mindless opening of mobile applications. We hypothesize that, to reduce the number of times that participants open the application, it would be more effective to, for example, create a prompt that, when participants try to open the application, checks that they really wish to do so. Although not in the scope of our experiment, we note that such alternative behavior change mechanisms provide rich opportunities for future research.

One year follow-up reveals diverse Facebook habits
Similar to Pielot and Rello [45] who followed up with participants several months after the end of an experiment on curbing digital overload, we followed up with our participants on their digital habits one year after the end of our study. Out of 50 participants, only 21 (15 males, 6 females) completed our follow-up survey on MTurk and nine of them reported that they currently used Facebook too much on their phone. However, all participants reported that they did not use any mobile application to assist them in regulating their digital consumption even though 19 out of 21 indicated that it is important for one to regulate time spent on apps. This finding suggests that although participants may consider their digital habits problematic they may likely not be ready to take actions to change their habit.

In place of using mobile apps to monitor phone usage habits, participants adopted different techniques such as using battery level: *"I check the battery usage for the usage of app"*; and level of fatigue: *"just an approximate calculation of how much my eyes get tired after using [my phone]"*. Half of the participants reported that in the last few months, they have encouraged others—such as friend, colleague, father, wife, brother-in-law, and sisters—to reduce how much time is spent on Facebook. It is likely that effects that participants consider beneficial may be passed on to others in their social circles. However, future work is needed to provide concrete evidence about these findings in our follow-up survey.

6 DISCUSSION

The high-level goal of our research was to create a generalizable technique that reduces digital consumption on mobile devices. To achieve this goal, we designed, implemented, and evaluated a mobile app intervention by combining nudge theory [53] from behavioral economics and negative reinforcement [51] from the theory of operant conditioning in behavioral psychology. In this section, we discuss how our intervention achieves behavior change, extends prior HCI research in digital consumption, and lessons learned from using MTurk to conduct a multi-week study.

6.1 Curbing Digital Overload

In contrast to HCI research that regulates technology overuse by blocking users' access to target services [23, 29, 33], our intervention allows users to continue accessing the target application when it is used past a daily limit. However, the user experiences a continual vibration nudge that creates an aversive environment for the participant. To avoid this negative stimulus, participants sensitive to the vibration stop using the target application after a short time. This unique combination of digital nudge and negative reinforcement provides evidence of an effective way to curb digital overload.

However, this change in behavior did not persist when the vibration intervention was withdrawn, which suggests that achieving long-term behavior change may require participants to experience negative reinforcement for an extended period of time (and perhaps indefinitely) or that the intervention should be applied during specific periods of time. In particular, the intervention could be used to achieve behavior change in specific contexts, such as when studying for an exam, during class time, when participating at a meeting that requires full focus, and more. In these kinds of scenarios, the vibration could serve as a reminder and help to nudge participants back to the task at hand.

Although we found that negative reinforcement via vibration was generally successful at reducing digital consumption of a target application, our attempts to personalize participants' usage limits did not reduce application usage significantly more than a rigid limit. This suggests a need for future experimentation that focuses on different personalization strategies. For example, the system could perform adaptive personalization that regularly recomputes usage limits over time instead of the static personalization used in our experiment. In addition, the system could be designed to integrate personal preferences, such as sleep time, wake-up time, or recreation time, as well as contextual awareness, such as location, time of day, weekday or weekend, vacation settings, and more. Determining how best to personalize or optimize the intervention so that it most effectively helps people to reduce their consumption of digital content is an exciting area for future work.

6.2 Extending Prior Research in Digital Consumption

There are a number of key differences between our work and recent research in HCI that focus on reducing digital consumption. For example, one major difference between our study and prior work ([16, 25]) is that our experiment did not explicitly target participants who self-identified as wanting to curb their digital consumption. We made this choice to reduce the effect of experiment confounds and because the concepts used to design our intervention do not require participants to be motivated for the intervention to be successful. Achieving significant behavior change without requiring participants to indicate their readiness to change makes our work to ecologically generalize to real-world settings where users do not often indicate their readiness to change. We hypothesize that future experiments conducted with only participants who are interested in reducing their digital consumption could result in an even larger effect.

Another key difference between our study and existing productivity tools ([41, 49, 54]) is that we provide automatic, non-invasive, realtime feedback that can be ignored or viewable by the user. By contrast, many existing productivity tools require users to visit a separate website or open a standalone application that displays their usage data and provides feedback. However, research has shown these tools suffer from low levels of engagement because of the need for users to make an effort to check a separate website or application [7]. Our solution provides a persistent yet non-intrusive notification that is always easily accessible and that displays realtime usage statistics that continuously update. In providing this continuous feedback, our work contributes to existing literature that aims to improve user engagement through visual feedback [10, 25, 50].

Finally, many prior solutions require that participants manually choose the applications that will be monitored, set their own personal goals regarding usage limits, and modify configurations ([16, 25]) but this burdens the users with more responsibility and might not be sustainable in the long-run. By contrast, our intervention is administered centrally and automatically monitors application usage without requiring the participant to take extra actions. We made this choice so that it was possible for us to conduct a controlled experiment in which all participants received exactly the same procedure. Although our approach relieves users of the burden of manually choosing the applications to monitor and deciding their own daily usage limits, individual participants will undoubtedly want to use and track different applications, as well as customize the intervention to their personal usage and habits. Providing a hybrid solution that removes a lot of user burden but provides enough control is an exciting direction for future research.

6.3 Lessons for future Mobile Studies on Crowdsourcing Platforms

In recent years, many behavioral science and HCI researchers have conducted experiments involving short tasks that engage workers recruited via crowdsourcing platforms for a few minutes or hours ([4, 30, 36]. But there exists ample opportunities to conduct mobile studies that run for several weeks with participants from crowdsourcing platforms and our experiment demonstrates a successful example with MTurk. These platforms are beneficial because they are cheaper, less time-consuming, and contain more diverse participants pool [48]. We discuss some of our lessons learned from this experimental context.

Throughout our recruitment cycle—from pilot studies to the final three-week experiment—we reached about 650 participants and eventually completed our experiment with 50 participants.

We conducted our experiment with only 10% of all participants reached because of our unique requirements: we first selected only users who used android phones, next we removed participants who did not already have Facebook mobile app installed on their phones before the study, then we removed participants who used Facebook for less than five minutes per day, and finally we selected MTurk workers located in India. Given these study constraints, it would have been challenging to recruit as many participants using traditional, non-online recruitment methods (such as university recruitment) as we would not have access to such a large participant population. Based on our experience, we provide the following recommendations for HCI researchers who want to explore the potential for using crowdsourcing platforms for mobile field studies:

Establish a live dashboard to monitor ongoing participation:
In contrast to short tasks that require minimal effort, setting up and monitoring our mobile study took a relatively longer time and effort. Since we wanted to continuously monitor participants' app usage habits beyond the first day of on-boarding, we created a live dashboard to keep track of the participants who enrolled in our study. The dashboard provided details about participants' phones to check for compatibility with our vibration tool, whether they had successfully enrolled in our experiment, whether their phones were actively used throughout the experiment, and a search functionality to find immediate information regarding specific participants. Without this dashboard, it would have been challenging to quickly filter out participants who enrolled in the study but who dropped out after the first few days. Using the dashboard, we were also able to generate a unique code for each participant to prevent scenarios where one participant might take part in the same experiment using multiple devices in an effort to cheat and receive multiple payment. These approaches point to the need for specialized tools that make it easy to conduct longitudinal experiments on MTurk or similar platforms.

Respond quickly to participants who are facing challenges: Effectively supporting crowdsourced workers throughout a multi-week study may be especially challenging due to the lack of opportunities to develop rapport with participants. As a result, it is crucial that researchers are available and highly responsive to participant inquiries as the study progresses. At the beginning of our experiment, we encouraged participants to reach out via MTurk if they had any questions. Then, to decrease the amount of time that it took us to respond to participants, we created a short script that extracted up-to-date information regarding each participant's ongoing enrollment from the dashboard and that also checked if the participant had completed all the necessary steps to enroll in the experiment. This enabled us to provide personalized responses that answered participant questions and addressed concerns in short periods of time.

Help participants to correctly follow instructions by integrating reminder features:
To ensure that participants followed the necessary steps to participant in our experiment, we implemented two features in our tool to assist participants at the beginning and end of the experiment. The first feature generated a unique code that each participant had to submit on MTurk. This code was generated only after a participant had successfully completed the enrollment steps. Using this approach enabled us to automatically filter out participants who started the process but did not pay attention to the enrollment instructions—these participants did not receive a code and were unable to continue with the experiment. At the end of the experiment, our application automatically enabled the second feature, which consisted of a notification reminder to participants to uninstall the application since the experiment was over. This is more effective than contacting participants through emails as participants may not check their emails regularly. Features like these helped us to spend more time successfully conducting our experiment and less time coordinating participants.

Recruit participants from specific locations to avoid confounds:
Our experiment was conducted with a diverse pool of participants, which is an advantage when compared to traditional non-online behavior change experiments that have been overwhelmingly conducted with participants from Western, Educated, Industrialized, Rich, and Democratic (WEIRD) societies [34]. At the beginning of our enrollment, we recruited participants from any locale, although we later changed our enrollment requirements to consist of only participants from India in order to avoid confounds that may arise due to cultural or geographic differences. However, we observe that the ability to recruit participants from specific countries could provide researchers with unique opportunities to investigate how their findings may differ based on participants recruited from different countries and cultures.

6.4 Generalizability and Limitations

We acknowledge that our study has a number of limitations. We do not apply advanced statistical models in our analysis and some findings were inconclusive due to low statistical power. Future work with a larger sample size and a longer study duration will address these issues. In addition, we only measured participants' application usage on their mobile phones, but it is possible that they may have used other electronic devices to access the application whenever their phones vibrated from overuse. Our work also experiments with only vibration, however, we hope to explore the dynamics of other forms of digital nudges in future work, such as screen dimming, phone LED blinking, WiFi disconnection, gentle sounds or ringtones, and more. Findings from experiments with different forms of digital nudge combined with negative reinforcement could provide a clearer picture of how to control digital distractions and advance research on human productivity, technology overuse, and digital addiction.

Our experiment focused only on a single target application: Facebook. Although we implemented our system to work on any mobile application, and tested it with a range of different applications, further experiments are necessary to confirm that our findings generalize beyond Facebook. In addition, instead of performing negative reinforcement on a per-app level, negative reinforcement could apply to the entire device, such as alerting a user when they are excessively checking their phone or when it is unlocked for long periods of time. The existence of such habit reinforcing systems could reduce the amount of cyberloafing that students engage in [55], help workers when they become distracted for too long [35], remind people to engage more in face-to-face interactions, and better balance people's digital consumption patterns.

7 CONCLUSION

This paper describes the design, implementation, and evaluation of a digital intervention that targets the reduction of digital overload. We design this intervention by uniquely combining the concepts of digital nudge and negative reinforcement to create a gentle vibration that nudges users to stop using a target mobile application after a daily usage limit is exceeded. We evaluated our approach through a three-week controlled experiment with 50 participants recruited from Amazon's MTurk platform. Our findings show that digital nudge successfully reduced digital overload by over 20% with statistically significant differences between study periods. Although the effect did not persist once the intervention was removed, qualitative feedback suggested that the system encouraged users to pay more attention to their personal usage habits and therefore that the intervention might be particularly effective when applied episodically and in specific settings such as education, sleep, and work. Our findings offer concrete benefits to the HCI community in the form of a theoretically grounded and easy-to-adopt component that could be integrated into broader interventions.

ACKNOWLEDGMENTS

The authors would like to thank Mashfiqui Rabbi and Jason Waterman for their feedback in this research project. The authors would also like to thank the reviewers for their valuable comments and helpful suggestions. The work is supported by funding from NSF (#134458) and Oath through the Connected Experiences Laboratory at Cornell Tech.

REFERENCES

[1] Amazon. 2017. Mechanical Turk. (2017). Retrieved Feb 1, 2018 from https://www.mturk.com.

[2] Morgan G Ames. 2013. Managing mobile multitasking: the culture of iPhones on stanford campus. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, 1487–1498.

[3] Eric PS Baumer, Phil Adams, Vera D Khovanskaya, Tony C Liao, Madeline E Smith, Victoria Schwanda Sosik, and Kaiton Williams. 2013. Limiting, leaving, and (re) lapsing: an exploration of facebook non-use practices and experiences. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 3257–3266.

[4] Michael Buhrmester, Tracy Kwang, and Samuel D Gosling. 2011. Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6, 1 (2011), 3–5.

[5] Rafael A Calvo and Dorian Peters. 2014. *Positive computing: technology for wellbeing and human potential.* MIT Press.

[6] Emily IM Collins, Anna L Cox, Jon Bird, and Cassie Cornish-Tresstail. 2014. Barriers to engagement with a personal informatics productivity tool. In *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: the Future of Design*. ACM, 370–379.

[7] Emily IM Collins, Anna L Cox, Jon Bird, and Daniel Harrison. 2014. Social networking use and RescueTime: the issue of engagement. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 687–690.

[8] Sunny Consolvo, Predrag Klasnja, David W McDonald, Daniel Avrahami, Jon Froehlich, Louis LeGrand, Ryan Libby, Keith Mosher, and James A Landay. 2008. Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays. In *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 54–63.

[9] Sunny Consolvo, David W McDonald, and James A Landay. 2009. Theory-driven design strategies for technologies that support behavior change in everyday life. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 405–414.

[10] Sunny Consolvo, David W McDonald, Tammy Toscos, Mike Y Chen, Jon Froehlich, Beverly Harrison, Predrag Klasnja, Anthony LaMarca, Louis LeGrand, Ryan Libby, and others. 2008. Activity sensing in the wild: a field trial of ubifit garden. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1797–1806.

[11] The Economist. 2017. Facebook, the world's most addictive drug. (2017). Retrieved May 13, 2017 from http://www.economist.com/blogs/graphicdetail/2016/04/daily-chart-5?fsrc=scn/fb/te/bl/ed/facebooktheworldsmostaddictivedrug.

[12] Ido Erev and Ernan Haruvy. 2016. Learning and the economics of small decisions. In *The handbook of experimental economics*, John H Kagel and Alvin E Roth (Eds.).

[13] Ido Erev and Alvin E Roth. 2014. Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences* 111, Supplement 3 (2014), 10818–10825.

[14] Deborah Estrin. 2014. Small data, where n= me. *Commun. ACM* 57, 4 (2014), 32–34.

[15] Rebecca Hansson, Peter Ljungstrand, and Johan Redström. 2001. Subtle and public notification cues for mobile devices. In *International Conference on Ubiquitous Computing*. Springer, 240–246.

[16] Alexis Hiniker, Sungsoo Ray Hong, Tadayoshi Kohno, and Julie A Kientz. 2016. MyTime: Designing and Evaluating an Intervention for Smartphone Non-Use. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4746–4757.

[17] Alexis Hiniker, Shwetak N Patel, Tadayoshi Kohno, and Julie A Kientz. 2016. Why would you do that? predicting the uses and gratifications behind smartphone-usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 634–645.

[18] Eva Hornecker, Stuart Swindells, and Mark Dunlop. 2011. A mobile guide for serendipitous exploration of cities. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. ACM, 557–562.

[19] Cheng-Kang Hsieh, Faisal Alquaddoomi, Fabian Okeke, John P Pollak, Lucky Gunasekara, and Deborah Estrin. 2018. Small Data: Applications and Architecture. In *Proceedings of the Fourth International Conference on Big Data, Small Data, Linked Data and Open Data*.

[20] Brian A Iwata. 1987. Negative reinforcement in applied behavior analysis: An emerging technology. *Journal of Applied Behavior Analysis* 20, 4 (1987), 361–378.

[21] Reynol Junco. 2013. Comparing actual and self-reported measures of Facebook use. *Computers in Human Behavior* 29, 3 (2013), 626–631.

[22] Reynol Junco and Shelia R Cotten. 2011. Perceived academic effects of instant messaging use. *Computers & Education* 56, 2 (2011), 370–378.

[23] Inyeop Kim, Gyuwon Jung, Hayoung Jung, Minsam Ko, and Uichin Lee. 2017. Letfis FOCUS: Mitigating Mobile Phone Use in College Classrooms. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 63.

[24] Jaejeung Kim, Chiwoo Cho, and Uichin Lee. 2017. Technology Supported Behavior Restriction for Mitigating Self-Interruptions in Multi-device Environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 64.

[25] Young-Ho Kim, Jae Ho Jeon, Eun Kyoung Choe, Bongshin Lee, K Kim, and Jinwook Seo. 2016. TimeAware: Leveraging framing effects to enhance personal productivity. In *Proceedings of the SIGCHI conference on human factors in computing systems*.

[26] Aniket Kittur, Ed H Chi, and Bongwon Suh. 2008. Crowdsourcing user studies with Mechanical Turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 453–456.

[27] Minsam Ko, Seungwoo Choi, Subin Yang, Joonwon Lee, and Uichin Lee. 2015. FamiLync: facilitating participatory parental mediation of adolescents' smartphone use. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 867–878.

[28] Minsam Ko, Seungwoo Choi, Koji Yatani, and Uichin Lee. 2016. Lock n'LoL: group-based limiting assistance app to mitigate smartphone distractions in group activities. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 998–1010.

[29] Minsam Ko, Subin Yang, Joonwon Lee, Christian Heizmann, Jinyoung Jeong, Uichin Lee, Daehee Shin, Koji Yatani, Junehwa Song, and Kyong-Mee Chung. 2015. NUGU: a group-based intervention app for improving self-regulation of limiting smartphone use. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. ACM, 1235–1245.

[30] Steven Komarov, Katharina Reinecke, and Krzysztof Z Gajos. 2013. Crowdsourcing performance evaluations of user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 207–216.

[31] Uichin Lee, Subin Yang, Minsam Ko, and Joonwon Lee. 2014. Supporting Temporary Non-Use of Smartphones. In *CHI14 Workshop: Refusing, Limitng, Departing: Why We Should Study Technology Non-Use Workshop*.

[32] LiveMint. 2018. India now has highest number of Facebook users, beats US. (2018). Retrieved Jan 23, 2018 from http://www.livemint.com/Consumer/CyEKdaltF64YycZsU72oEK/Indians-largest-audience-country-for-Facebook-Report.html.

[33] Markus Löchtefeld, Matthias Böhmer, and Lyubomir Ganev. 2013. AppDetox: helping users with mobile app addiction. In *Proceedings of the 12th international conference on mobile and ubiquitous multimedia*. ACM, 43.

[34] Edouard Machery. 2010. Explaining why experimental behavior varies across cultures: a missing step in fithe weirdest people in the world?fi. *Behavioral and brain sciences* 33, 2-3 (2010), 101–102.

[35] Gloria Mark, Shamsi T. Iqbal, Mary Czerwinski, and Paul Johns. 2015. Focused, Aroused, but so Distractible: Temporal Perspectives on Multitasking and Communications. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 903–916.

[36] Winter Mason and Siddharth Suri. 2012. Conducting behavioral research on Amazonfis Mechanical Turk. *Behavior research methods* 44, 1 (2012), 1–23.

[37] Melissa Mazmanian. 2013. Avoiding the trap of constant connectivity: When congruent frames allow for heterogeneous practices. *Academy of Management Journal* 56, 5 (2013), 1225–1250.

[38] Brandon T McDaniel and Sarah M Coyne. 2016. "Technoference": The interference of technology in couple relationships and implications for womenfis personal and relational well-being. *Psychology of Popular Media Culture* 5, 1 (2016), 85.

[39] Adrian Meier, Leonard Reinecke, and Christine E Meltzer. 2016. "Facebocrastination"? Predictors of using Facebook for procrastination and its effects on studentsfi well-being. *Computers in Human Behavior* 64 (2016), 65–76.

[40] Nicholas Micallef, Lynne Baillie, and Stephen Uzor. 2016. Time to exercise!: an aide-memoire stroke app for post-stroke arm rehabilitation. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. ACM, 112–123.

[41] Moment. 2017. Automatically track your and your family's daily iPhone and iPad use. (2017). Retrieved April 2, 2017 from https://inthemoment.io/.

[42] Fabian Okeke, Michael Sobolev, and Deborah Estrin. 2018. Towards A Framework for Mobile Behavior Change Research. In *Proceedings of the Technology, Mind, and Society*. ACM, 27.

[43] Martin T Orne. 1962. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist* 17, 11 (1962), 776.

[44] Martin Pielot and Rodrigo de Oliveira. 2013. Peripheral vibro-tactile displays. In *Proceedings of the 15th international conference on Human-computer interaction with mobile devices and services*. ACM, 1–10.

[45] Martin Pielot and Luz Rello. 2017. Productive, Anxious, Lonely: 24 Hours Without Push Notifications. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, Article 11, 11 pages. DOI:http://dx.doi.org/10.1145/3098279.3098526

[46] Benjamin Poppinga, Charlotte Magnusson, Martin Pielot, and Kirsten Rassmus-Gröhn. 2011. TouchOver map: audio-tactile exploration of interactive maps. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. ACM, 545–550.

[47] Andrew K Przybylski, Kou Murayama, Cody R DeHaan, and Valerie Gladwell. 2013. Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior* 29, 4 (2013), 1841–1848.

[48] David G Rand. 2012. The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of theoretical biology* 299 (2012), 172–179.

[49] RescueTime. 2016. Find your ideal work-life balance. (2016). Retrieved September 13, 2016 from https://www.rescuetime.com.

[50] John Rooksby, Parvin Asadzadeh, Mattias Rost, Alistair Morrison, and Matthew Chalmers. 2016. Personal tracking of screen time on digital devices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 284–296.

[51] Burrhus Frederic Skinner. 1953. *Science and human behavior*. Simon and Schuster.

[52] Matthew J Stollak, Amy Vandenberg, Andie Burklund, and Stephanie Weiss. 2011. Getting social: The impact of social networking usage on grades among college students. In *Proceedings from ASBBS annual conference*, Vol. 18. 859–865.

[53] Richard H Thaler and Cass R Sunstein. 1999. *Nudge: Improving decisions about health, wealth, and happiness*. HeinOnline.

[54] unGlue. 2017. unGlue: End screen time battles. Start better habits. (2017). Retrieved April 2, 2017 from http://www.unglue.com/.

[55] Yiran Wang, Melissa Niiya, Gloria Mark, Stephanie M. Reich, and Mark Warschauer. 2015. Coming of Age (Digitally): An Ecological View of Social Media Use Among College Students. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 571–582. DOI:http://dx.doi.org/10.1145/2675133.2675271