

App Usage Predicts Cognitive Ability in Older Adults

Mitchell L. Gordon*
Stanford University
mgord@cs.stanford.edu

Leon Gatys
Apple Inc
lgatys@apple.com

Carlos Guestrin
Apple Inc
guestrin@apple.com

Jeffrey P. Bigham
Apple Inc
jbigham@apple.com

Andrew Trister
Apple Inc
atrister@apple.com

Kayur Patel
Apple Inc
kayur@apple.com

ABSTRACT

We have limited understanding of how older adults use smartphones, how their usage differs from younger users, and the causes for those differences. As a result, researchers and developers may miss promising opportunities to support older adults or offer solutions to unimportant problems. To characterize smartphone usage among older adults, we collected iPhone usage data from 84 healthy older adults over three months. We find that older adults use fewer apps, take longer to complete tasks, and send fewer messages. We use cognitive test results from these same older adults to then show that up to 79% of these differences can be explained by cognitive decline, and that we can predict cognitive test performance from smartphone usage with 83% ROCAUC. While older adults differ from younger adults in app usage behavior, the “cognitively young” older adults use smartphones much like their younger counterparts. Our study suggests that to better support all older adults, researchers and developers should consider the full spectrum of cognitive function.

ACM Reference Format:

Mitchell L. Gordon, Leon Gatys, Carlos Guestrin, Jeffrey P. Bigham, Andrew Trister, and Kayur Patel. 2019. App Usage Predicts Cognitive Ability in Older Adults. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300398>

1 INTRODUCTION

Although more than 42% of older adults use smartphones [3], relatively little is known about smartphone usage patterns amongst older adults (65 or older), particularly when

*This work was done while Mitchell L. Gordon was an intern at Apple Inc.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5970-2/19/05.

<https://doi.org/10.1145/3290605.3300398>

compared to what we know about the younger population’s smartphone usage patterns [27]. Understanding similarities and differences in how older adults and the younger population use smartphones is important for designing engaging and helpful experiences. Emerging trends in health and wellness applications delivered through smartphones seem especially well-suited to the needs of older adults, yet designers may fail to create applications that address the needs of older adults without this understanding. For example, older adults often experience cognitive decline while aging, but we do not know to what extent this affects smartphone use [26].

In this paper, we present a quantitative study of smartphone usage by older adults (aged 61 to 76), utilizing a dataset that we created containing usage metadata logs of 84 iPhone users over a three month period. We characterize application usage in a number of different ways, including by type and number of applications used, amount of time spent on applications, and behavior in switching between applications. We compare with prior work that looked at similar behavior in the general population, noting that the usage behaviors of older adults are different than younger adults, but more similar than might have been previously assumed. This suggests that approaches developed for the general population might usefully be applied for older smartphone users as well.

We also uncover important differences that may guide design principles for application developers to better design for the older population. By reproducing analyses frequently performed on other smartphone usage data sets in prior work, we demonstrate that older adults display distinct app usage characteristics from younger adults, such as using fewer apps and keeping them open for longer. To further characterize the usage patterns of older adults, we apply new analysis techniques for app-usage dynamics to reveal how applications are used in sequence (*e.g.*, most switches into Photos comes from Camera, and most switches into Safari come from Mail).

While these differences likely result from a complex combination of generational and age-related shifts in culture, values, lifestyle, and cognitive / physical abilities, we focus on how cognitive function of participants is associated with differences in use. Cognitive function is of particular interest

in the context of older adults, as cognitive function has previously been found to be important across a wide array of activities, *e.g.*, in technology usage [16], office productivity [29], and avoiding motor vehicle accidents [55].

Each participant in our study also underwent neurocognitive testing in two major cognitive domains: (i) perceptual and motor, and (ii) learning and memory. To explore how differences in cognitive function might explain differences we observe in behavior, we constructed a quasi-experiment implemented using our usage log data paired with the neurocognitive testing results of our participants. Our 84 participants had assessed neurocognitive function in the normal range, yet we observed a wide distribution of performance on neurocognitive testing. Specifically, some healthy older adults perform on neurocognitive testing at the level of younger adults, while others did not. We compared app usage behavior of those who performed like younger adults with the remainder of our participants, and found that the observed differences in usage behavior can be largely explained by cognitive function. We generalized these results into a predictive model capable of predicting neurocognitive testing performance from app usage patterns with 83% ROC AUC.

Our research demonstrates that while older adults display some notable differences as compared to younger adults, “cognitively young” older adults use smartphones much like their younger counterparts. To better support all older adults, researchers and developers should especially consider how to design applications that can accommodate users who are no longer cognitively young. Older adults now form a sizable percentage of the user base of smartphones. Our research contributes an important step toward understanding how to best support them.

2 PRIOR WORK

Our work is informed by prior work on (i) studies of smartphone app usage patterns, (ii) how older adults use technology, and (iii) predicting personal characteristics from smartphone usage.

Smartphone App Usage Patterns

Smartphone app usage patterns have been studied extensively. Researchers have investigated how many apps people install, which apps are most used at which time of day, and how long these apps are kept open for [10, 19, 24]. They’ve investigated daily routines [20] and even predicted which app a user will launch next [5].

We note that two factors unify most of this work. The first is that these studies have typically relied upon data from younger adults. For instance: in Do *et al.*’s study, using data gathered from 77 volunteers, only 11% of participants were over 40 [19]. Larger-scale studies, which rely on users

deciding to download publicly available logging apps, typically did not collect demographic data but can be assumed to contain primarily younger adults given the demographics of Android users [10]. While some specific demographics have been investigated, such as a 14-participant study of teenagers [9], researchers have yet to characterize the usage patterns of older adults.

The second unifying factor is that, as pointed out by Morrison *et al.* [36], nearly all of this work has been performed using Android phones rather than on iPhone and iOS, as the Android platform allows for the easy installation of logging software for OS-level events such as app launches and phone unlocking. Morrison responded by reproducing on iPhone a variety of studies previously performed on Android. Morrison achieved this by focusing on 10,000 jailbroken users, which enabled them to install an app that could log OS-level events. While this study did not collect or report demographic information, people who jailbreak their iPhones have been found to have a mean of 36 years old and are 65% male [35].

In this paper, we analyze iPhone usage patterns – notably, older adults are more likely to use an iPhone than any other age groups [22]. Therefore, when we compare usage statistics reproduced on our population with other studies, we do so primarily with Morrison *et al.*’s work, given that it is the only study reporting general usage patterns from iPhones. iOS is different than Android and may promote different usage patterns. Further, people who choose iPhone have been found to exhibit sociological differences [43]. One difficulty with this comparison is that, as noted by Morrison, users of jailbroken devices likely differ from the general population in that they may be more engaged and familiar with their iPhones. Therefore, their usage patterns can likely be considered on the more ‘extreme’ end of what we may expect from the general population [36].

Technology and Older Adults

Older adult’s relationship with technology is an important area of research and has been studied extensively. Technology use tends to decline with age [46]. While older adults generally use fewer technologies than younger adults, and use them less frequently, notable exceptions have been found in domains such as healthcare [11, 40].

While older adult’s general relationship with technology is well-documented, their specific usage patterns are not. Work which quantifies how older adults use smartphones – increasingly the most important and versatile pieces of technology in our lives – is extremely limited. Many studies have broadly surveyed older adults on their preferences and attitudes towards mobile phones, but have not analyzed actual usage data [38]. Some studies are centered around input. For instance, older adults have been found to input

text via voice equally as fast as younger adults, but slower in other input methods [49]. While some work has sought to create guidelines for designing mobile phone interfaces for older adults, typically recommending larger buttons and less complex interfaces [1, 18, 47], it relies primarily on surveys of older adults and their preferences, rather than analysis of their actual app usage data.

Even more difficult than understanding the app usage patterns of older adults is understanding *why* these patterns might be different than younger adults. One area of study is cognitive ability. Older adults with cognitive difficulties may perceive technology as more difficult to use [45], and individuals with higher cognitive levels are likely to participate in a wider variety of web-based tasks (i.e., Internet use and email) [23]. Perceptual speed, an area of cognitive ability, has been found to be a reasonable predictor of determining desktop browsing characteristics in older adults [15]. Potential cultural and lifestyle causes have also been explored. In a qualitative survey of 14 older adults, several reasons for a lack of technology use emerged: they were found to be uncomfortable with having to take on responsibility for tasks previously handled by trained professionals, made deliberate decisions to not use technologies when they perceived them as replacing or eroding something of value to them; and supported by the fact that opting out of technology use fits with cultural expectations and thus seems acceptable, despite being increasingly limiting in digital society [31]. In our work, we aim to quantify how much of certain app-usage behavioral differences between older adults and younger adults can be explained by cognitive ability versus other factors.

Relating Personal Characteristics to Smartphone Use

Smartphones can be a window into their users' minds. Early work in this area found that smartphone usage patterns can be used to predict personality traits [12]. More recently, smartphone usage patterns have been used to predict propensity to trust others, their college GPA, depression dynamics, sleep quality, and social capital [6, 8, 48, 51, 52]. Researchers have even shown that people exhibit unique app usage signatures, and can be differentiated just from their usage patterns [56]. Other work has focused on linking real-time physiological traits to usage patterns. In a study of 20 participants over 40 days, Murnane *et al.* connected app usage patterns with biological rhythms, finding that usage patterns vary for individuals with different body clock types [37].

To our knowledge, our work is the first to focus on how smartphone app usage patterns can be explained by cognitive ability. Other work has looked at how higher-level smartphone habits may have a lasting impact on cognitive function and development [54]. We are aware of only one paper which predicts cognitive ability from smartphone usage, and its focus was on lower level physical touch events such

Dataset statistics	
# of participants	84
Age range	61-76
Median age	66
% female	69.0
Observation period	Three months, mid 2018
Total app launches	494,641
Total unlocks	186,968
Total minutes of app use	1,396,429.9
Total messages	386,573
Total calls	66,371

Table 1: Descriptive statistics from our dataset of smartphone usage log data from older adults.

as tapping on the space-bar followed by the first character of a word, rather than on higher level app usage behavior [17].

3 DATASET DESCRIPTION

We conducted a three-month observational study using smartphone data collected from 84 healthy older adults, ranging in age from 61 to 76, with paired neurocognitive testing. Our dataset contains 494,641 app launches and 186,968 phone unlocks. Table 1 lists descriptive statistics from our dataset.

All participants were existing iPhone users (iPhone 5S or newer) for at least one year prior to joining the study. Participants were provided with an iPhone 7 Plus to use as their primary personal phone. The iPhones were instrumented with logging software that recorded lock/unlock events, apps launched/closed, and anonymized metadata for each text message and call sent/received. A full backup of their personal pre-study iPhone was uploaded to the new iPhones 7. We gave participants new iPhones installed with a backup of their existing iPhone for two reasons: first, logging app launches on iOS requires custom software that could not be installed on participants' existing phones. Second, doing so removes confounders related to individual device quality, speed, size, or battery.

At the start of the study, participants took a neurocognitive screening exam called the Cogstate Brief Battery (CBB). We discuss the CBB in more detail in Section 5. No participants with diagnosed cognitive impairment were included in our study. Participants were recruited at medical centers in several locations around the United States. Participants completed a survey to report diagnosed medical conditions. We removed participants diagnosed with cognitive impairment, deafness, blindness, motor impairment disorders, or depression. Participants were not given special instruction as to how frequently or for what purpose they should use their iPhones. We surveyed participants for their race, educational attainment, and employment status. 82 participants identified themselves as White Caucasian, 1 as African American, and 1 as Hispanic. 34% had achieved a graduate degree

as their highest level of education, 35% an undergraduate degree, and 31% a high school diploma. 26% were currently employed full time, 13% part-time, 53% retired, and 6% other.

For the purposes of our evaluation, we consider an app launch to take place each time a user actively brings a new app into focus. This includes when a user taps on a notification to launch its corresponding app, but does not include notifications themselves, widgets, or other background services. Following convention from prior work, when app categories are discussed, we use the categories provided by app developers to the Apple App Store. For Apple’s first party apps that are built into iOS and are thus not on the Store, we manually assigned apps to categories.

4 HOW OLDER ADULTS USE SMARTPHONES

We begin by reproducing analyses by Morrison *et al.* [36]. As mentioned in our Prior Work section, this is the only other paper to have characterized iPhone app-usage statistics. Their analyses consist primarily of descriptive statistics centered around number of apps used, type of apps used, total usage by time of day, and usage duration. In the following subsections, we report our results and discuss how they compare. In many cases, the scope of our analysis is limited by available prior work. In many cases, we go beyond what has been done in prior work in order to present a more thorough picture of how older adults use smartphones, despite the inability to directly compare those results to younger adults. Thus, our work combines replication of prior work with new analysis, focusing on older adults whose usage has not been studied in this way before.

Number of apps and duration of use

Participants used an average of 50.3 unique apps (std. dev 22.7). This is notably lower than Morrison *et al.*, who found an average of 64.0 (std. dev. 52.5) and the Boehmer *et al.* Android study which found an average of 74.37 (std. dev. 44.16) [10]. We find that older adults keep apps open for an average of 121.2 seconds, nearly 37% longer than the 88.6 seconds found by Morrison *et al.*

App usage by category

In Table 2, we aggregate app usage by category. We list the percent of app launches that fall into each category along with the mean duration of app usage. In a comparison to the results reported Morrison *et al.*’s work (who also used the App Store’s categories), we find that the Communication category dominates for both populations, though with older adults at a slightly lower 46.3% than Morrison’s 49%. However, Social Networking apps make up only half as many of launches for older adults as they do for younger adults, at 6.8% compared to 12.6%. When older adults do launch a Social Networking app, they keep it open for nearly twice as long, 207 vs. 115

Category	Avg. duration	Pct. launches	Sample apps
Communication	74.2	46.32	Messages, Mail, Phone, Messenger
Browser	144.7	8.32	Safari, Google Chrome
Productivity	70.3	6.80	Gmail, Calendar, Notes, Reminders
Photo & Video	103.1	6.55	Instagram, Camera, Photos, YouTube
Social Networking	207.3	6.37	Facebook, LinkedIn, Nextdoor, Pinterest
Utilities	107.1	4.63	Google, Clock, App Store, Calculator
Games	364.4	4.04	Words With Friends, Candy Crush
Navigation	226.5	2.59	Google Maps, Apple Maps, Waze
News	196.6	2.48	Apple News, BuzzFeed, NY Times
Health & Fitness	67.5	1.73	Activity, FitBit, Health, MyFitnessPal
Settings	68.1	1.61	Settings
Weather	53.6	1.28	Weather, The Weather Channel
Finance	55.7	1.19	E*TRADE Mobile, Chase Mobile, Wallet, Coinbase
Music	159.0	1.19	Apple Music, Pandora, Spotify
Travel	73.2	0.98	Yelp, Uber, Lyft, Expedia
Business	94.2	0.81	Slack, Kisi
Shopping	116.2	0.69	Amazon, Poshmark, Groupon
Books	168.6	0.59	Audible, Kindle, eBooks
Entertainment	183.2	0.56	Podcasts, Netflix, iTunes Store, Hulu
Lifestyle	96.0	0.53	Redfin Real Estate, Nest, Zillow
Food & Drink	111.7	0.27	Starbucks, OpenTable, Caviar
Sports	238.8	0.21	ESPN, FanDuel, 18 Birdies Golf
Reference	237.9	0.11	Bible, Dictionary
Education	194.6	0.08	Duolingo, Lumosity: Daily Brain Games, iTunes U
Medical	100.2	0.07	Mychart, Mango Health, Migraine Buddy

Table 2: Average duration and percent launches for each category, sorted by most launched. Categories for each app are taken from the App Store.

seconds. Perhaps the biggest difference in a single category is Weather: the category makes up only 0.3% of launches in younger adults, but is more than quadrupled for older adults at 1.3%. Duration of Weather apps is significantly longer for older adults as well, at 53.6 seconds vs. 28.9 seconds.

App usage by hour of day

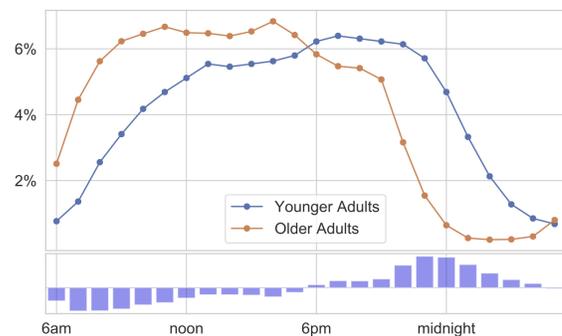


Figure 1: Older adults use their phones earlier in the day than younger adults. We show app launches by hour of day from our dataset of older adults overlaid on the results reported by Morrison et al [36] for younger adults. The y-axis is normalized as percent of total usage to allow for comparison between populations with different sample sizes.

Figure 1 shows the total number of app launches per hour of the day, overlaid atop the same plot from Morrison et

al’s work among younger users. We note a distinct shift in older adult’s usage to earlier in the day. This may be suggestive of different sleep patterns between the populations, as smartphone usage has been found to provide a reasonable measure of sleeping habits [14].

App micro-usage

Ferreira *et al.* introduced the concept of a *micro-usage* [21], which defines short bursts of interaction with apps versus longer, more involved interactions. Typically, micro-usage is found by partitioning app usage duration data into two clusters, which look for a natural break point that can separate a micro-usage from a longer usage. To compute this break-point, we follow prior work by applying k-means clustering to log-scale app duration data, with $k=2$. The micro-usage breakpoint we found was 27.5 seconds, 28.5% higher than the 21.4 seconds found by Morrison *et al.* We propose two possible explanations for this difference: that older adults attempt to complete more tasks during micro-uses of apps, or that older adults take longer to complete the same tasks during micro-uses as younger adults. Unfortunately, we can’t conclusively investigate this with our data, but it may be an interesting question for future work.

Multi-app sessions

Following Morrison *et al.*, a session is defined as the period of activity between unlocking and locking a phone, with the requirement that a session is ended only after the phone has been locked for at least 30 seconds. We found that participants open an average of 2.9 apps per session, with 67.2% of sessions consisting of 3 apps or fewer, and 73.5% of sessions consisting of 4 apps or fewer. 36.1% of sessions contain only one app. By comparison, Morrison *et al.* found 76.7% containing 3 apps or fewer, 83.5% containing 4 apps or fewer, and 38% containing only one app. These results suggest that older adults open more apps per session than younger adults. Morrison does not report average number of apps per session.

App dynamics

Inspired by work relating to behavioral patterns during web surfing sessions [53], we believe that deeper exploration of sequential app usage within and between sessions is an interesting and important component in understanding the app-usage behavior of users. Specifically, we focused on behavior around how users switch from one app to another. Though not studied in great detail in prior literature, some work has used sequential app usage patterns as a means to predict or recommend future apps [33]. We aim to understand and visualize these patterns, and call this area of study “app dynamics”. In this paper, we limited our scope of exploration to sequential app usage within single sessions. To

analyze this data, we created bi-grams from each sequence. For instance, if within a single session a user opened Facebook, followed by Messages, followed by Weather, we would then have the bi-grams (Facebook, Messages) and (Messages, Weather). We then treat these bi-grams as a bipartite graph, where apps being switched out of and apps being switched into are treated as disjoint sets of nodes, with edges indicating a switch between two apps. We find 289,771 resulting switches in our dataset. We chose to remove all self-switches from our analysis. While likely an interesting and worthy area of study, self-switches are complicated by unclear intent (*e.g.*, someone may have accidentally closed an app or be returning from a phone call). We believe that extracting meaning from these switches would require further logging of OS-level behaviors.

Figure 2 shows the frequency and direction of app switches for the top twelve apps in our dataset (in terms of total launches), created from the bi-grams discussed above. The thickness of the lines indicate magnitude in terms of percent of switches. While for clarity of visualization we only plot the top twelve apps, we note that 72% of sessions start with one of the top twelve apps, and 70.7% of switches are into one of the top 12 apps.

We note a few interesting trends. A sizable percent of switches from every app are into Messages. Given that Messages displays a notification when a new message is received, and that incoming messages are often of high interest to smartphone users [50], this behavior makes sense – users stop what they are doing when they receive a message to attend to that message. On the opposite end of the spectrum, we see that Contacts is rarely switched into, even from Phone. We observe that Messenger (Facebook’s messaging product) shows a strong link to Facebook, while Facebook has a much wider distribution of potential switches. Future work, particularly with a larger sample size, should work to understand and visualize how these switches may differ between users.

Switch duration. Above, we studied the origin and destination of app switches. We now investigate the time it takes for those switches to occur. Specifically, we are interested in switches where the user chooses to move directly from one app to another app, without engaging in other iOS-related features such as widgets along the way, which adds noise to the process of switching. For this analysis, we therefore filtered to only switches where there are two sequential apps logged as coming into focus with no other process coming into focus between them (the home screen itself is not logged by the OS as coming into focus when it appears). To further reduce the chance that a user is not directly switching from one app to another (as opposed closing an app, putting their phone down for a little while without locking it, and picking

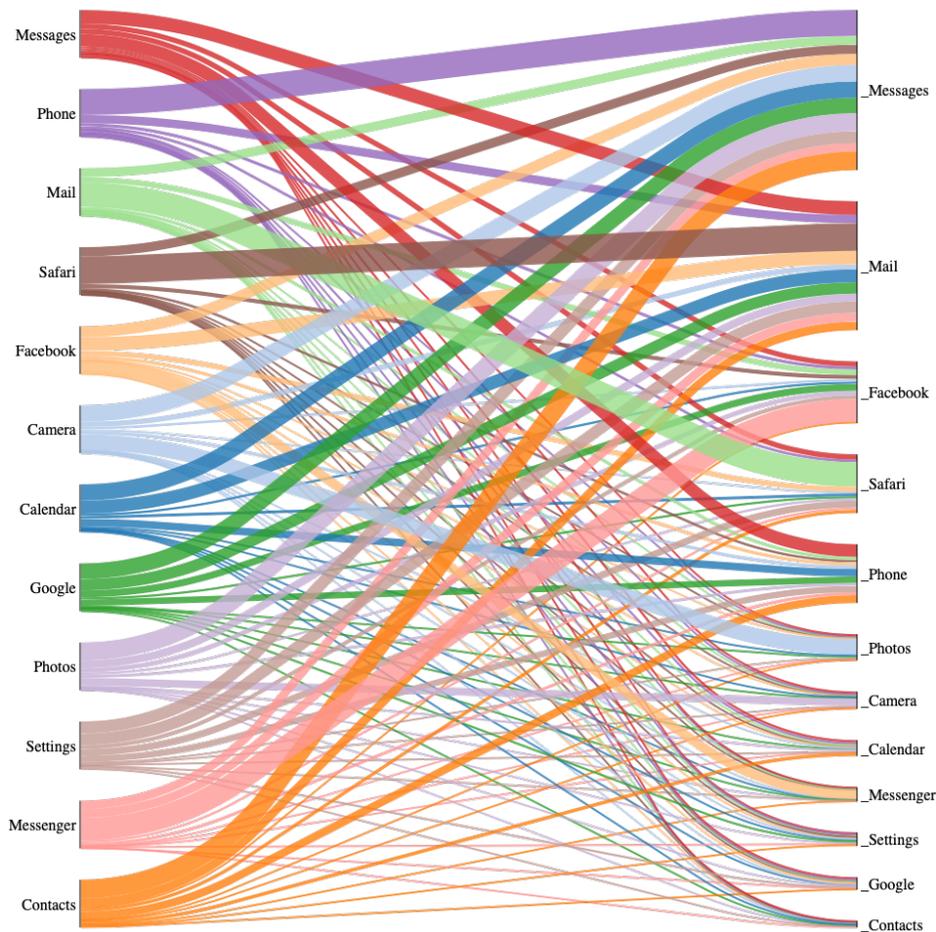


Figure 2: Frequency and direction of app switches, as percent of total switches, for the top twelve apps in our dataset. Messages is the most frequently switched into app, making up 31% of switches. We see clear connections between some apps, such as Camera to Photos and Messenger to Facebook.

it back up again), we filtered to switches that took at most 30 seconds. This left us with 120,721 switches. We found that a switch takes an average of 10.79 seconds.

Messages and calls

We found that our users send and receive an average of 37.2 messages and 5.1 calls per day. While we are not aware of any iPhone-specific studies reporting the messaging and calling behavior of younger adults, a platform-agnostic Pew study found that teenagers send and receive a mean of 67 text messages per day [32].

5 EXPLAINING BEHAVIORAL DIFFERENCES THROUGH COGNITIVE ABILITY

We provide evidence above that older adults display notably different app usage patterns than younger adults. We now ask: how can these differences be explained? As mentioned

earlier, in this paper we focused specifically on cognitive function. We aim to quantify how much of the difference between usage patterns of younger adults and older adults can be explained by cognitive decline. To do this, we defined a group of *cognitively young* older adults and investigated how these users differ from younger adults, as well as from cognitively normal adults.

Cogstate Brief Battery Assessment

To answer this question, we make use of our app usage data in combination with cognitive test results from each participant.

During study enrollment, the CogState Brief Battery assessment (CBB) was administered to each participant. The CBB consists of a simple reaction time task (Detection task), a choice reaction time task (Identification task), a one-back working memory task (One-Back task), and a continuous

recognition visual learning task (Learning Task). A full description of each task and their validity is available in Maruff *et al.* [34]. For each individual task, Cogstate reports many associated raw performance measures, often involving accuracy and speed of performance. Depending on the task, one of the raw scores is designated by Cogstate as the “primary outcome” of the task. Cogstate also provides two composite scores computed from a mixture of these raw scores and normalized across ages. The two composites are described as follows: Psychomotor Function/Attention composite, incorporating the Detection and Identification tests, relative to age-matched normative data, and Learning/Working Memory composite, incorporating the One Card Learning and One Back tests, relative to age-matched normative data. Within our population of healthy older adults, we find that our CBB scores are approximately normally distributed.

Experimental design

To determine whether differences in app usage patterns between younger adults and older adults can be explained by cognitive function, we designed a quasi-experiment. Quasi-experiments are frequently used to estimate causal relationships in health-related observational studies where randomized controlled studies would otherwise be infeasible, unethical, or impossible [28]. In our case, random assignment would be unethical and likely impossible since it would require recruiting a group of older adults with equal cognitive ability and then intentionally inducing cognitive decline to a random subset of those participants.

For our quasi-experiment, we employed a *posttest-only nonequivalent control group design* [13]. We created two groups. The first group consisted of older adults with the cognitive ability of a typical younger adult. We denote this group as YC. To determine which participants to assign to this group, we relied on normative data provided by Cogstate on their CBB test, which indicates the scores expected of a healthy younger adult. We selected all older adults in our study whose CBB scores are within one standard deviation of the normative data, and assign them to the YC group. We then created a matched group, NC, from the remaining participants, using an almost-exact matching strategy [44], matching on the age covariate. This is important because age is known to be correlated with cognitive decline and, as we demonstrated earlier, app usage patterns. We created these groups for the two different composites mentioned above. After creating our groups in this way, we had 16 participants per group.

Results

For each metric reported in the “How Older Adults Use Smartphones” section, we report the mean value from our YC and

NC groups along with statistical significance from a Mann-Whitney U test. We also report Pearson correlation values for each metric continuously over the entire dataset (not just the matched users in the YC and NC groups). Table 3 shows the results from our quasi-experiment. For each app usage metric, to test whether our result may be confounded by socio-economic factors, we also ran an ANCOVA analysis to look for a main effect of race, educational attainment, or employment status on app usage patterns, after controlling for cognitive scores. We do not find significant main effects, meaning we don’t find evidence these covariates are confounding the results of our quasi-experiment.

Quasi-experiment Discussion

The results of our quasi-experiment indicate that a surprisingly large amount of the difference in app usage behavior between older adults and younger adults can be explained by cognitive ability. In this section, we discuss the results from each app-usage behavior tested in the experiment. We found that cognitive ability can explain over 67% of the difference between older adults and younger adults in terms of number of apps used, over 76% of the difference in duration, and 79% of the difference in micro-usage breakpoint.

Number of apps used. We found that number of apps used varies significantly with the Psychomotor Function/Attention composite, but not with Learning/Working Memory composite. YCs in Psychomotor Function/Attention use 59.5 apps on average, while NCs use only 47.9. Compared to older adults’ overall mean of 50.3 reported earlier, Psychomotor Function/Attention can explain over 67% of the difference from the 64.0 found for Morrison *et al.*’s younger adults. This result agrees with findings from psychological literature: we hypothesize that number of apps used may be a proxy for information seeking and curiosity, which has been found to be strongly linked with attention [25].

App duration. Duration varies significantly with the Learning/Working Memory composite, with YCs averaging 96.2 seconds compared to NCs 122.4. Learning/Working Memory can explain over 76% of the difference between the duration for younger adults, which Morrison *et al.* found to be 88.6 seconds, and older adults overall which we found to be 121.2 seconds. We hypothesize that how long a user keeps an app open may be a proxy for processing speed, and prior work has found that working memory influences processing speed and reading fluency [30]. Therefore, we expect that YCs for this composite would experience this faster usage duration.

Micro-usage breakpoint. The micro-usage breakpoint for YCs in the Learning/Working Memory composite is 7.2 seconds lower. We did not find a significant difference for the Psychomotor Function/Attention composite. Compared to our

Behavior	Learning/Working Memory			Psychomotor Function/Attention		
	YC mean	NC mean	Pearson R	YC mean	NC mean	Pearson R
# apps used	47.8	51.0	.038	59.5*	47.9*	.256*
App duration (seconds)	96.2**	122.4**	-.220*	119.6	116.3	.033
Micro-usage breakpoint (seconds)	22.7***	29.9***	-.256*	28.5	28.9	.039
Apps per session	2.6	3.1	-.032	3.2	2.9	.095
Messages per day	43.7*	33.8*	.143	36.5	37.6	0.070
Calls per day	9.9*	6.4*	.189	7.1	7.9	-.092
App switch speed	10.09*	10.74*	-.279**	10.56	10.62	-.099
Early morning pct usage	12.54	10.1	0.228*	11.46	10.38	0.27*
Mid morning pct usage	29.56**	26.54**	0.136	27.88	26.98	0.162
Afternoon pct usage	33.44	33.91	-.033	34.77	33.57	-.008
Night pct usage	22.08**	27.12**	-0.346**	24.46	26.51	-0.253*

Table 3: Results from our quasi-experiment. Statistically significant results are highlighted in bold. Older adults who are cognitively young in Learning/Working Memory keep apps open for less time and use their phones earlier in the morning. Older adults who are cognitively young in Psychomotor Function/Attention use more apps. Mann-Whitney U test between YC and NC columns is used to test significance (* $p < 0.05$, ** $p < 0.01$, * $p < 0.001$).**

overall older adults micro-usage breakpoint of 27.5 seconds, as reported earlier, cognitive ability can explain 79% of the difference from the 21.4 second breakpoint in younger adults reported by Morrison *et al.* [36].

Categorical/apps breakdown. As in the prior section of this paper, we investigated both duration and percent of total usage in terms of categories. We also analyzed potential differences in specific apps (only the top twenty apps in terms of total launches were selected). Given there are 20 apps and 25 categories, each with two dimensions to evaluate them on, there is not room to list each result. Rather, we chose to list none of them, as no significant correlations were found after correcting for multiple hypothesis testing (even before correction, only three significant correlations were found). One likely explanation is that as we evaluated granular behaviors in our relatively small dataset, noise between participants became too great to statistically overcome.

Messages. We found that the number of messages varies significantly with Learning/Working Memory, with YCs sending and receiving 10 messages more per day. As mentioned earlier, we are not aware of comparable data from younger iPhone users. This result is in agreement with psychology literature, which has found that social engagement has significant links with cognitive ability among older adults [7].

Switching speed. We found that YCs in the Learning/Working Memory composite switch apps 2/3 of a second faster than NCs. One might expect to also see a difference in the Psychomotor Function composite, given that switching apps quickly likely requires some level of quick physical tapping coordination. Surprisingly, we do not see a significant result here, indicating that the likely reason older adults switch

apps more slowly is due to their mental capacity to retain the information necessary to complete the tasks, rather than the physical ability to tap the correct areas.

Hour of day. Here we analyze the percent of app launches that fall into a given time of day. Daytime sleepiness has been found to be highly correlated with cognitive ability in older adults [39], and we hypothesized that daytime sleepiness may manifest itself as relatively fewer launches during the morning and daytime versus night. We broke down app usage by time into four components: early morning (5am-7am), mid-morning (8am-12pm), afternoon (12pm-5pm), and night (5pm-10pm). Our results appear to confirm this relationship, particularly when looking at learning/working memory. While 22.1% of usage takes place at night for YCs, 27.1% takes place at night for NCs. Interestingly, this is our only result in which YCs actually behave *less* like younger users: earlier, we observed that younger adults tend to use their phones far more during late night hours than older adults. Figure 3 shows this difference.

6 PREDICTING COGNITIVE ABILITY

We designed a novel prediction task of predicting “cognitively young” older adults – predicting whether an older adult has the cognitive ability of a typical younger adult. To build such a model, we leveraged insights described in previous sections. We demonstrated that despite the large variation inherent to human behavior, and the relatively small sample size in our dataset, we can achieve up to 83% ROC AUC for learning/working memory. This prediction task is designed to validate our empirical findings, and show that the results we report above can be used as generalizable features in a predictive model.

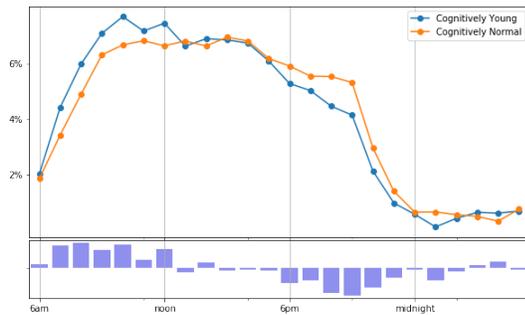


Figure 3: Cognitively young older adults launch a larger percentage of apps in the morning, and cognitively normal older adults launch a larger percentage at night.

Task Description

Based on a user’s app-usage behavior, we predicted whether that user has the cognitive ability of a typical younger adult or has experienced cognitive decline. We attempt this task for both of the CBB composites described in the previous section, and using the same normative data described in the previous section to determine which users fall into the cognitively young or cognitively normal groups.

Experimental setup

For this classification task, we used the area under the ROC curve (ROC AUC) as our evaluation metric with Leave One Out cross-validation for estimation. We experiment with L1-normed logistic regression and Gradient Boosted Trees [42], finding that the latter produced better results. (71% vs. 83% ROC AUC). To avoid over-fitting, for models where the number of feature vectors exceeds ten, we select the ten best features using the training set of each Leave-One-Out iteration, calculated from their ANOVA F-values. The hyper-parameter max tree depth is determined by a 20-fold cross-validation using the training set within each LOO iteration. As age-corrected normative data were used to compute the CBB composites (*i.e.*, Psychomotor/Attention and Learning/Working Memory), age was not entered as a covariate in these analyses. Because of the unbalanced dataset (20.5% of participants were cognitively young) and the trade-off between true-and false-positive rate associated with prediction, we chose to compare models using the area under the receiver operating characteristic (ROC) curve (AUC) which is equal to the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one [2]. Thus, a random baseline will score 50%.

Models. To illustrate the predictive power of the different types of app-usage behaviors reported earlier in this paper, we define a series of models, each with a different feature set corresponding to one of these app-usage behaviors:

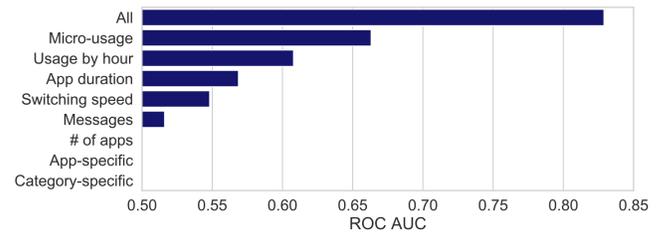


Figure 4: Accuracy for predicting cognitively young older adults in terms of the Learning/Working Memory composite for each feature set. We reach 83% ROC AUC in a combined Gradient Boosted Trees model with all features. For the Psychomotor Function/Attention composite, only one feature predicted better than random guessing: number of apps used reached 72% ROC AUC.

- (1) **Number of apps used:** Number of unique apps opened.
- (2) **App duration:** Mean duration across all apps opened.
- (3) **Micro-usage:** A user’s micro-usage breakpoint and percent of app launches that are a micro-usage.
- (4) **Category-specific:** For each category, percent launches and mean duration of an app use in that category.
- (5) **App-specific:** For each app, percent launches and mean duration of an app use for that app.
- (6) **Messages:** Mean number of messages per day.
- (7) **Switching speed:** Average time taken to switch between apps (switches filtered as described earlier).
- (8) **Usage by hour of day:** Percent app launches by hour of the day.
- (9) **All:** All features combined.

Results

Learning/Working Memory. Figure 4 shows the results from predicting whether an older adult is cognitively young in terms of learning/working memory. With a Gradient Boosted Trees model that chooses from all available features, we reach 83% ROC AUC. We note that micro-usage breakpoint, app switch duration, and usage by time of day result in the highest prediction performance. The combined model outperforms each of the feature sets individually, suggesting that together the features complement each other with new information, rather than capturing the same information.

Psychomotor Function/Attention. Predicting older adults who are cognitively young in terms of Psychomotor function / attention from our app-usage behaviors proved to be significantly more challenging. We find that only one feature produced an accuracy better than random guessing: number of apps used, which produced an accuracy of 72% ROC AUC. This is not surprising, as we found few significant results for this composite in our earlier quasi-experiment, but number

of apps used showed a strong result with cognitively young older adults using 11.6 more apps than cognitively normal older adults.

7 DISCUSSION

There are several important limitations to our work. The first stems from our sample size. At 84 participants, our dataset is larger than many papers which discuss app usage patterns and which connect smartphone usage to psychological or behavioral characteristics. However, an even larger dataset would enable analysis of more granular behaviors such as specific app usage patterns and also would have enabled to better control for confounders. While we made a strong effort to account for confounders that seemed most likely to impact our results, future work should further investigate these. For instance, while we remove participants with cognitive impairment, deafness, blindness, motor impairment disorders, and depression, there may be potential other medical confounders that we were unable to consider. A larger and more diverse sample size would also enable further investigation of potential socio-economic confounders.

Another important limitation is that we did not directly recruit younger adults to compare older adults with; rather, we compare our results with prior work published on younger adults. This resulted in two challenges: first, we were only able to compare results for metrics that we replicate from prior work (rather than those we created ourselves). Second, while we took care to reproduce the studies we compare our results to as closely as possible, reproduction always opens up the door to potential differences in data collection, cleaning, or analysis methods that may influence results. We also note that, by giving users new iPhones 7 installed with a backup of their existing iPhones, it is possible that some participants could experience a breaking-in period to adjust to their new physical devices which could potentially impact how they used their phones.

The final limitation relates to our quasi-experiment. While quasi-experiments are a common and generally accepted method to estimate causal direction in behavioral health research [28], we note that they are still estimates and not guaranteed to be correct to the same degree a controlled randomly-assigned experiment would be.

Our results underscore prior qualitative results that older adults are not a monolithic population. We believe there may be opportunities to create accessibility features for smartphones that can take in to account the spectrum of cognitive ability and its impact on usage patterns. We find that while cognitively young older adults are able to hold more information and be more expressive in their usage patterns and choices (1, 3, 4), and therefore have similar engagement patterns when they do use phones, they are still effected by the physical results of aging (2). Therefore, our results show that

a sizable number of older adults – those who are cognitively young – may be best treated as similar to a younger population for features relating only to cognitive ability. However, as we see when looking at usage by time of day, older adults who are cognitively young do not necessarily display signs of physical/lifestyle youth, and, in fact, moved in the opposite direction from how younger adults behaved (Table 3).

There are two important implications of our work. First, our results suggest that smartphones could be used as passive preliminary detection tools for cognitive decline by monitoring app usage patterns. This potentially enables alerts to loved ones or suggestions that users seek professional medical advice. Second, prior work has described methods to design interfaces while minimizing cognitive load [41]. Related methods have been used to design smartphone interfaces intended specifically for older adults [4]. However, these are likely not equally desirable for all older adults because they can come at the cost of efficiency or utility. Our results suggest that smartphones and their apps may be able to dynamically adjust their interfaces to better support the individual cognitive needs of older adults by passively detecting cognitive decline.

8 CONCLUSION

We present a quantitative analysis of the smartphone app usage patterns of older adults. We find that older adults use fewer apps, keep them open longer, and use their phones earlier in the day. We also introduce analysis on what we call “app usage dynamics”, in which we investigate which apps users switch into and how long it takes them to perform the switch. We then design a quasi-experiment to explain the observed differences between older adults and younger adults through cognitive ability, finding that usage patterns of “cognitively young” older adults come close to matching those of younger adults for some metrics like number of apps used, but not others like the times of day they use their phones. We then show that the results presented in our quasi-experiment are sufficiently strong to predict older adults who are cognitively young with 83% ROC AUC. While prior work suggests that differences in older adults usage of smartphones may be primarily attributable to differences in culture, values, and lifestyle, we find that for some behaviors, differences can be largely explained by cognitive decline.

ACKNOWLEDGEMENTS

We thank our colleagues at Apple, including Gustav Larsson, Nick Foti, Myoung Cha, Brian Kent, Jim Kretlow, Sherry Wu, Silvana Ciurea-Ilcus, and Younghoon Kim for their contributions.

REFERENCES

- [1] Muna S Al-Razgan, Hend S Al-Khalifa, Mona D Al-Shahrani, and Heshah H AlAjmi. 2012. Touch-based mobile phone interface guidelines and design recommendations for elderly people: A survey of the literature. In *International Conference on Neural Information Processing*. Springer, 568–574.
- [2] Tim Althoff and Jure Leskovec. 2015. Donor retention in online crowdfunding communities: A case study of donorschoose.org. In *Proceedings of the 24th international conference on world wide web*. International World Wide Web Conferences Steering Committee, 34–44.
- [3] MONICA Anderson and ANDREW Perrin. 2017. Technology use among seniors. *Washington, DC: Pew Research Center for Internet & Technology* (2017).
- [4] Farah Arab, Yasir Malik, and Bessam Abdulrazak. 2013. Evaluation of PhonAge: an adapted smartphone interface for elderly people. In *IFIP Conference on Human-Computer Interaction*. Springer, 547–554.
- [5] Ricardo Baeza-Yates, Di Jiang, Fabrizio Silvestri, and Beverly Harrison. 2015. Predicting the next app that you are going to use. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM, 285–294.
- [6] Yin Bai, Bin Xu, Yuanchao Ma, Guodong Sun, and Yu Zhao. 2012. Will you have a good sleep tonight?: sleep quality prediction with mobile phone. In *Proceedings of the 7th International Conference on Body Area Networks*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 124–130.
- [7] Shari S Bassuk, Thomas A Glass, and Lisa F Berkman. 1999. Social disengagement and incident cognitive decline in community-dwelling elderly persons. *Annals of internal medicine* 131, 3 (1999), 165–173.
- [8] Ghassan F Bati and Vivek K Singh. 2018. Trust Us: Mobile Phone Use Patterns Can Predict Individual Trust Propensity. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 330.
- [9] Frank Bentley, Karen Church, Beverly Harrison, Kent Lyons, and Matthew Rafalow. 2015. Three hours a day: understanding current teen practices of smartphone application use. *arXiv preprint arXiv:1510.05192* (2015).
- [10] Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger, and Gernot Bauer. 2011. Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage. In *Proceedings of the 13th international conference on Human computer interaction with mobile devices and services*. ACM, 47–56.
- [11] Jennifer K Carroll, Anne Moorhead, Raymond Bond, William G LeBlanc, Robert J Petrella, and Kevin Fiscella. 2017. Who uses mobile phone health apps and does use matter? A secondary data analytics approach. *Journal of medical Internet research* 19, 4 (2017).
- [12] Gokul Chittaranjan, Jan Blom, and Daniel Gatica-Perez. 2013. Mining large-scale smartphone data for personality studies. *Personal and Ubiquitous Computing* 17, 3 (2013), 433–450.
- [13] Larry B Christensen, Burke Johnson, Lisa Anne Turner, and Larry B Christensen. 2011. Research methods, design, and analysis. (2011).
- [14] Matthew A Christensen, Laura Bettencourt, Leanne Kaye, Sai T Moturu, Kaylin T Nguyen, Jeffrey E Olgin, Mark J Pletcher, and Gregory M Marcus. 2016. Direct measurements of smartphone screen-time: relationships with demographics and sleep. *PLoS one* 11, 11 (2016), e0165331.
- [15] Michael Crabb and Vicki L Hanson. 2016. An analysis of age, technology usage, and cognitive characteristics within information retrieval tasks. *ACM Transactions on Accessible Computing (TACCESS)* 8, 3 (2016), 10.
- [16] Sara J Czaja, Neil Charness, Arthur D Fisk, Christopher Hertzog, Sankaran N Nair, Wendy A Rogers, and Joseph Sharit. 2006. Factors predicting the use of technology: findings from the Center for Research and Education on Aging and Technology Enhancement (CRE-ATE). *Psychology and aging* 21, 2 (2006), 333.
- [17] Paul Dagum. 2018. Digital biomarkers of cognitive function. *npj Digital Medicine* 1, 1 (2018), 10.
- [18] Ana Correia de Barros, Roxanne Leitão, and Jorge Ribeiro. 2014. Design and evaluation of a mobile user interface for older adults: navigation, interaction and visual design recommendations. *Procedia Computer Science* 27 (2014), 369–378.
- [19] Trinh Minh Tri Do, Jan Blom, and Daniel Gatica-Perez. 2011. Smartphone usage in the wild: a large-scale analysis of applications and context. In *Proceedings of the 13th international conference on multimodal interfaces*. ACM, 353–360.
- [20] Katayoun Farrahi and Daniel Gatica-Perez. 2008. Daily routine classification from mobile phone data. In *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 173–184.
- [21] Denzil Ferreira, Jorge Goncalves, Vassilis Kostakos, Louise Barkhuus, and Anind K Dey. 2014. Contextual experience sampling of mobile application micro-usage. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. ACM, 91–100.
- [22] Fluent. 2016. Devices & Demographics 2016.
- [23] Jeremy Freese, Salvador Rivas, and Eszter Hargittai. 2006. Cognitive ability and Internet use among older adults. *Poetics* 34, 4-5 (2006), 236–249.
- [24] Andrea Girardello and Florian Michahelles. 2010. AppAware: Which mobile applications are hot?. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*. ACM, 431–434.
- [25] Jacqueline Gottlieb, Pierre-Yves Oudeyer, Manuel Lopes, and Adrien Baranes. 2013. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences* 17, 11 (2013), 585–593.
- [26] Caroline N Harada, Marissa C Natelson Love, and Kristen L Triebel. 2013. Normal cognitive aging. *Clinics in geriatric medicine* 29, 4 (2013), 737–752.
- [27] Eszter Hargittai, Anne Marie Piper, and Meredith Ringel Morris. [n. d.]. From internet access to internet skills: digital inequality among older adults. *Universal Access in the Information Society* ([n. d.]), 1–10.
- [28] Anthony D Harris, Jessina C McGregor, Eli N Perencevich, Jon P Furuno, Jingkun Zhu, Dan E Peterson, and Joseph Finkelstein. 2006. The use and interpretation of quasi-experimental studies in medical informatics. *Journal of the American Medical Informatics Association* 13, 1 (2006), 16–23.
- [29] James J Heckman, Jora Stixrud, and Sergio Urzua. 2006. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor economics* 24, 3 (2006), 411–482.
- [30] Lisa A Jacobson, Matthew Ryan, Rebecca B Martin, Joshua Ewen, Stewart H Mostofsky, Martha B Denckla, and E Mark Mahone. 2011. Working memory influences processing speed and reading fluency in ADHD. *Child Neuropsychology* 17, 3 (2011), 209–224.
- [31] Bran Knowles and Vicki L Hanson. 2018. The wisdom of older technology (non) users. *Commun. ACM* 61, 3 (2018), 72–77.
- [32] Lenhart, Amanda, Pew Research Center. 2015. *Teen, Social Media and Technology Overview 2015*. [Online; accessed 20-September-2018].
- [33] Eric Hsueh-Chan Lu, Yi-Wei Lin, and Jing-Bin Ciou. 2014. Mining mobile application sequential patterns for usage prediction. In *Granular Computing (GrC), 2014 IEEE International Conference on*. IEEE, 185–190.
- [34] Paul Maruff, Elizabeth Thomas, Lucette Cysique, Bruce Brew, Alex Collie, Peter Snyder, and Robert H Pietrzak. 2009. Validity of the CogState brief battery: relationship to standardized tests and sensitivity

- to cognitive impairment in mild traumatic brain injury, schizophrenia, and AIDS dementia complex. *Archives of Clinical Neuropsychology* 24, 2 (2009), 165–178.
- [35] Donald McMillan, Alistair Morrison, and Matthew Chalmers. 2013. A Comparison of Distribution Channels for Large-Scale Deployments of iOS Applications. In *Developments in Technologies for Human-Centric Mobile Computing and Applications*. IGI Global, 222–239.
- [36] Alistair Morrison, Xiaoyu Xiong, Matthew Higgs, Marek Bell, and Matthew Chalmers. 2018. A Large-Scale Study of iPhone App Launch Behaviour. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 344.
- [37] Elizabeth L Murnane, Saeed Abdullah, Mark Matthews, Matthew Kay, Julie A Kientz, Tanzeem Choudhury, Geri Gay, and Dan Cosley. 2016. Mobile manifestations of alertness: Connecting biological rhythms with patterns of smartphone app use. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. ACM, 465–477.
- [38] Ofcom. 2018. Adults' Media Use And Attitudes Report.
- [39] Maurice M Ohayon and Marie-Françoise Vecchierini. 2002. Daytime sleepiness and cognitive impairment in the elderly population. *Archives of internal medicine* 162, 2 (2002), 201–208.
- [40] Katherine E Olson, Marita A O'Connell, Wendy A Rogers, and Neil Charness. 2011. Diffusion of technology: frequency of use for younger and older adults. *Ageing international* 36, 1 (2011), 123–145.
- [41] Sharon Oviatt. 2006. Human-centered design meets cognitive load theory: designing interfaces that help people think. In *Proceedings of the 14th ACM international conference on Multimedia*. ACM, 871–880.
- [42] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [43] Ahmad Rahmati, Chad Tossell, Clayton Shepard, Philip Kortum, and Lin Zhong. 2012. Exploring iPhone usage: the influence of socioeconomic differences on smartphone adoption, usage and usability. In *Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services*. ACM, 11–20.
- [44] Paul R Rosenbaum. 2002. Observational studies. In *Observational studies*. Springer, 1–17.
- [45] Lena Rosenberg, Anders Kottorp, Bengt Winblad, and Louise Nygård. 2009. Perceived difficulty in everyday technology use among older adults with or without cognitive deficits. *Scandinavian journal of occupational therapy* 16, 4 (2009), 216–226.
- [46] Neil Selwyn. 2004. The information aged: A qualitative study of older adults' use of information and communications technology. *Journal of Aging studies* 18, 4 (2004), 369–384.
- [47] Paula Alexandra Silva, Kelly Holden, and Philipp Jordan. 2015. Towards a list of heuristics to evaluate smartphone apps targeted at older adults: a study with apps that aim at promoting health and well-being. In *System Sciences (HICSS), 2015 48th Hawaii International Conference on*. IEEE, 3237–3246.
- [48] Vivek K. Singh and Isha Ghosh. 2017. Inferring Individual Social Capital Automatically via Phone Logs. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 95 (Dec. 2017), 12 pages. <https://doi.org/10.1145/3134730>
- [49] Amanda L Smith and Barbara S Chaparro. 2015. Smartphone text input method performance, usability, and preference with younger and older adults. *Human factors* 57, 6 (2015), 1015–1028.
- [50] Cary Stothart, Ainsley Mitchum, and Courtney Yehnert. 2015. The attentional cost of receiving a cell phone notification. *Journal of experimental psychology: human perception and performance* 41, 4 (2015), 893.
- [51] Rui Wang, Gabriella Harari, Peilin Hao, Xia Zhou, and Andrew T Campbell. 2015. SmartGPA: how smartphones can assess and predict academic performance of college students. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. ACM, 295–306.
- [52] Rui Wang, Weichen Wang, Alex daSilva, Jeremy F Huckins, William M Kelley, Todd F Heatherton, and Andrew T Campbell. 2018. Tracking Depression Dynamics in College Students Using Mobile Phone and Wearable Sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 43.
- [53] Weinan Wang and Osmar R Zaiane. 2002. Clustering web sessions by sequence alignment. In *Database and Expert Systems Applications, 2002. Proceedings. 13th International Workshop on*. IEEE, 394–398.
- [54] Henry H Wilmer, Lauren E Sherman, and Jason M Chein. 2017. Smartphones and cognition: A review of research exploring the links between mobile technology habits and cognitive functioning. *Frontiers in psychology* 8 (2017), 605.
- [55] Ides Y Wong, Simon S Smith, and Karen A Sullivan. 2012. The relationship between cognitive ability, insight and self-regulatory behaviors: Findings from the older driver population. *Accident Analysis & Prevention* 49 (2012), 316–321.
- [56] Sha Zhao, Julian Ramos, Jianrong Tao, Ziwen Jiang, Shijian Li, Zhaohui Wu, Gang Pan, and Anind K Dey. 2016. Discovering different kinds of smartphone users through their application usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 498–509.