How do people use their smartphone?

A data scientific approach to describe and identify user-related, system-related and context-related patterns in use

Andrew Hendrickson[1], PhD, Lieven De Marez[2], PhD, Marijn Martens[2] (MA), Gytha Muller[1], Koen Ponnet[2], PhD, Catherine Schweitzer[1], and Mariek M. P. Vanden Abeele[3], PhD

[1]Department of Cognitive Science & Artificial Intelligence, Tilburg University, the Netherlands

[2]Department of Communication Sciences, Ghent University, Belgium

[3]Department of Cognition and Communication, Tilburg University, the Netherlands

Abstract

Quantifying and understanding the myriad ways people use their phones and how that impacts their relationships, cognitive abilities, mental health, and well being is increasingly important in our phone-centric society. However, most studies on the patterns of phone use have focused on theory-driven tests of specific usage hypotheses using self-report questionnaires or analyses of smaller datasets. In this work we present a series of analyses from a large corpus of over 3000 users that combine data-driven and theory-driven analyses to identify reliable smartphone usage patterns and clusters of similar users. Furthermore, we compare the stability of user clusters across user- and system-initiated sessions, as well as during the hypothesized ritualized behavior times directly before and after sleeping. Our results indicate support for some hypothesized usage patterns but present a more complete and nuanced view of how people use smartphones.

Introduction

To date, no technology has diffused into society more rapidly and globally than the smartphone (ITU, 2017; Poushter & Stewart, 2016). The success of the smartphone can be attributed to its substantial contribution to our daily autonomy. Smartphones enable us to perform our social roles, manage our social networks and access personalized information and services independent from time and place constraints (Vanden Abeele, De Wolf, & Ling, 2018). While this ubiquitous connectivity brings new opportunities, it also poses new risks and challenges. Smartphone use is, for example, linked to mental health problems (e.g., Jenaro, Flores, Gómez-Vela, González-Gil, & Caballo, 2009; Twenge, Joiner, Rogers, & Martin, 2018), stress (Hsiao, Shu, & Huang, 2017), reduced productivity (e.g., David, Kim, Brickman, Ran, & Curtis, 2015; Spira & Feintuch, 2005; Vanden Abeele, Schouten, et al., 2016), poor sleep quality (Lanaj, Johnson, & Barnes, 2014), conflict in personal relationships (Vanden Abeele, Antheunis, & Schouten, 2016), poor work-life balance (Wajcman, Bittman, & Brown, 2008) and risky behaviors such as texting while driving (Bayer & Campbell, 2012). Given the promises and perils of smartphone use, it is not surprising that media effects researchers have embraced the study of the effects of smartphone use.

A majority of the extant research on the implications of smartphone use has used survey designs to identify how people's smartphone usage is related to their thoughts, emotions and behaviors. These survey studies typically involve theory-driven tests of specific usage hypotheses, using questionnaires in which users self-report on their smartphone use. These studies have been valuable; they have identified relationships between smartphone use and a myriad of antecedents and outcomes, thus contributing to theory formation and to societal valorization. Despite their value, however, they have important limitations.

A first limitation is that studies may present inaccurate or even invalid findings by including self-report measures of smartphone use that are known to be unreliable in their analysis; there is a large discrepancy between how frequently people say they use their smartphone, and what they actually do. The reason why self-report measures of smartphone use are unreliable is that the fragmented, yet frequent occurrence of the behavior is difficult for individuals to recall accurately. Hence, individuals instead rely on heuristic estimation strategies when self-reporting their phone use and these estimation strategies are biased. For example, research shows that heavy phone users underestimate their use, whereas light phone users overestimate their use (Vanden Abeele, Beullens, & Roe, 2013).

A second limitation of relying on self-reported smartphone use is that several studies limit themselves to operationalizations of smartphone use in terms of self-reported frequency and/or duration of use (e.g., Bayer & Campbell, 2012). However, smartphone use is fragmented and patterned behavior (Deng et al., 2018; Oulasvirta, Rattenbury, Ma, & Raita, 2011). It is this very nature of smartphone use that may produce unique effects; we should thus explore which patterns there are, and in which way they are meaningful. Scholars currently circumvent this issue by asking individuals to self-report about their patterned behavior (e.g., to indicate how habitual their phone use is), but we lack insight into the reliability of these assessments; and there is a risk that factors such as social desirability bias these responses.

Third, a limitation of survey designs is that they typically focus on smartphone use as a stable behavior that only shows variability across individuals, as within-person variability is difficult to assess via cross-sectional self-report measurements; this is unfortunate, as such intra-person variability may be meaningful. For example, specific patterns of smartphone behavior may vary depending on person-specific factors such as a person's mood or cognitive state. Moreover, self-reported smartphone use measures fail to account for variability that is system-specific or context-specific. With respect to system-specific variability, smartphone use may be associated with app- or device settings, for example, smartphone behaviors that manifest themselves in response to incoming notifications. With respect to context-specificity, smartphone use may be associated with a myriad of situational characteristics, such as those that pertain to the spatial, temporal or social environment in which the smartphone use takes place. For example, smartphone checking behaviors may manifest differently in situational contexts where the behavior is deemed inappropriate (e.g., while driving or during a romantic dinner) than where it is deemed appropriate (e.g., during waiting situations). However, most of the extant empirical work overlooks differences within the individual, which may, in turn, result in part from differences in device-generated events (e.g., did usage start in response to a notification or not) or differences in contexts (e.g., the time or place of use).

In recent years, scholars have begun to addresses the above shortcomings by exploring the use of smartphone logging data rather than self-reports in empirical studies. In addition to capturing what people do rather than what they say, smartphone logging data reveal the dynamic nature of smartphone use, in the form of within-person variability in smartphone use. Moreover, the log data can include system-generated smartphone cues and can contextualize the smartphone use if indicators of the spatio-temporal context are registered. Gathering log data results in a rich collection of time-dependent signals that can be made sense of by relying on computational research methods (e.g., Deng et al., 2018; Gouin-Vallerand & Mezghani, 2014; Hao, Wang & Xu, 2016). The use of computational methods to study the effects of smartphone use is promising, as it creates opportunities to explore complex system- and context-specific patterns of smartphone usage, both at the aggregate level, and at the level of one unique individual.

To date, studies presenting a computational analysis of smartphone log data are scarce in the field of communication sciences, and particularly that of smartphone effects research. Most of the extant work has focused on research questions in the area of product design (e.g., Wu, Liang & Tang, 2017) or computer science (e.g., Gouin-Vallerand & Mezghani, 2014). Moreover, these studies mostly rely on datasets garnered among relatively small sets of individuals (e.g., Hao, Wang & Xu, 2016). As a result, many meaningful questions about patterns in smartphone use and their social implications remain unexplored. This is unfortunate, as these patterns lie at the heart of research questions such as how fragmentation in smartphone use is linked to attention in various spheres of our everyday life (e.g., Williams, 2018). A reason why the use of smartphone log data is rare in the field of communication research may be that we currently have only a limited understanding of what patterns could potentially be distilled from log data using computational methods. Additionally, scholars may benefit from gaining access to low-threshold operational tools to subtract such patterns from log data. Hence, the purpose of this study is to provide an analysis of three patterns that can be distilled by drawing sets of data-driven features from smartphone log data. The patterns that we aim to identify are (1) app repertoire, (2) habitual smartphone use, and (3) fragmentation.

The app repertoire concept can be understood as an element of niche theory (Dimmick, Kline, & Stafford, 2016; Ramirez, Dimmick, Feaster, & Lin, 2008). The central assumption of niche theory is that media need to offer unique benefits to their users in order to survive. If media do not succeed in gratifying unique user needs, in other words, if media lack to create a unique niche for themselves, they will have to compete against other media who fulfil similar functions. As a result of media convergence, smartphone device give access to a large number of mobile applications that vary in functionality, and thus in the needs that they gratify. Using the niche theory lens, we may assume that some of these apps compete against each other, whereas others don't – for example, when a person uses Google as a search engine, it is unlikely that this same user also actively uses Bing, or another search engine, although these mobile applications may be installed on the phone.

Research shows that while people have several mobile applications downloaded and installed on their smartphone, they only use only a handful on a frequent basis (Jung, Kim & Chan-Olmsted, 2014) – this usage concentration in a few applications aligns with the notion of 'media repertoires' that users build, presumably to be better able to cope with complex media environments (Heeter, 1985 as cited in Jung, Kim & Chan-Olmsted, 2014, p. 354). Media repertoire can be explored both at the aggregate level, and at the level of unique individuals. At the aggregate level, information regarding the mobile applications in which usage is concentrated in *across* individuals, can inform about the niche that smartphones as media devices serve. Relying on smartphone log data, Jung, Kim and Chan-Olmsted (2014) found, for example, that usage is concentrated mostly in messaging and social media applications, partly due the simple fact that these apps have an 'inherent network externality value' (p. 356), meaning that there is a critical mass of users who use these networked apps to fulfil social, utility and belongingness needs, which makes the value of being part of the network greater relative to the cost of missing out. With respect to our study, this means that we may expect a similar aggregate pattern to reveal itself. In a second step, however, this aggregated information on usage concentration can be contrasted against individual app repertoires. Here, we may see that there may be inter-individual differences in the mobile apps in which users concentrate their use.

The app repertoire concept is related to the notion of habitual smartphone use. The smartphone is considered a device that embeds a reward infrastructure in the form of dynamic social and informational mobile applications such as social media and news apps that offer users access to a constant stream of stimulating social and informational content (Oulasvirta et al., 2011). In combination with system features (e.g., push notifications) and haptic feedback features (e.g., pull down menus) these dynamic content applications invite users to repeatedly engage with the smartphone and its contents, and thus support the development of habitual smartphone checking behaviors in individuals (Bayer, Campbell, & Ling, 2016; Oulasvirta et al., 2011). These checking behaviors can be linked to the notion of app repertoire, as habitual checking routines likely involve the selection and patterned use of those apps in which individuals concentrate their usage (Jung, Kim and Chan-Olmsted, 2014). Important to note, here, is that certain cues can serve as 'gateways' that activate the checking habit (Bayer, Campbell, & Ling, 2016) - with notifications being a central type of cue that may bring users to start checking their smartphone applications, but also other (contextual or within-person) cues may trigger checking habits. Additionally, scholars have described that dynamic content applications may elicit checking behavior whereby users access the same apps over and over again in one or multiple temporally close smartphone sessions or whereby a brief checking event serves as the gateway to a lengthy usage session (Bayer et al., 2016; Schnauber-Stockmann, Meier, & Reinecke, 2018), phenomena referred to as smartphone users ‘going down the rabbit hole’ (e.g., Collier, 2016)

Finally, a third pattern that is meaningful to explore is fragmentation. Smartphone use is highly dispersed behavior, intersecting with people’s activities from the early morning to their bedtime in the evening (Deng et al., 2018; Oulasvirta, Rattenbury, Ma, & Raita, 2011). This is in part an artefact of the use of messaging and social media applications. Such applications have radically altered our everyday practices by supporting near-synchronous communication: Mobile messaging interactions are essentially a-synchronous in that senders know that recipients may not read their message immediately, and thus do not have to worry about interrupting them, while receivers can choose an appropriate time to read their messages, after which they have time to edit their responses (Rettie, 2009). However, there are clear normative expectations to respond shortly after having read a message – these expectations are supported through features such as read-receipts (e.g., Whatsapp’s blue checkmarks) or status features (e.g., “is currently online”, “is typing”) (Ling & Lai, 2016). The near-synchronous nature of mobile communication has led to a ‘connected’ mode of relational management (Licoppe, 2004). Social media applications operate under a similar logic of near-synchronicity by offering users extensive notification-features that inform them of every new like, comment, or update from people within their network.

Given that smartphone users who use messaging and social media applications are in a constant state of connected presence (cf. Licoppe, 2004), we may expect that there is high fragmentation in these users' smartphone sessions across the day; the fragmented use of messaging and social media applications may trigger users to habitually check other apps high in their respective app repertoire.

App repertoire, habitual checking behavior and fragmentation are patterns that have been conceptually identified and explored in extant studies on the social implications of smartphone use. It is meaningful, however, to empirically explore these patterns further. A data-driven approach is most valuable for such an exploration, as it unveils how these patterns take form in reality, how these patterns are clustered, and how stable they are both within and across individuals. Moreover, a data-driven approach enables to assess how stable these patterns are when accounting for system-generated cues, such as notifications, and when accounting for context, such as at what times these patterns take place throughout the day. The variability in this person-, device- and context-specificity can be operationalized in subsequent studies testing predictive relationships between smartphone use patterns and certain antecedents or outcomes.

Exploring these questions is crucial if we wish to advance our understanding of the complex and patterned nature of smartphone use, and how its person-, device- and context-specific nature may produce unique effects on aspects such as users' health and wellbeing. For example, if we consider the relationship between smartphone use and work-life balance, it is most relevant to explore inter-individual differences in the position of work email in individual's app repertoires, whether checking work email is part of users' habitual checking routines, to what extent checking work email is fragmented during the day, and whether notifications from the email app and the daytime influences the aforementioned patterns – the same user may experience work-life balance differently, for example, on days where s/he wakes up with new email notifications and checks work email before bedtime than on days where work email notifications are disabled and work emails is only checked during working hours.

The aims of the study are threefold: First, we aim to identify three patterns that are conceptually meaningful: app repertoire, habitual checking behavior and fragmentation. This aim is achieved via an aggregated analysis of features that can be drawn from smartphone log data. Second, we explore to what extent there is inter-individual variability in the examined patterns, and how they themselves cluster together. This aim is achieved by the clustering of individuals on the basis of the derived features, which also informs about which features are key to drive the clustering process. Third, we aim to explore the person-, system-, and context-specificity of these patterns. This aim is achieved by testing stability of the patterns within individuals, by itself, and when accounting for notifications and the time of day.

Method

In this section we first outline how data was collected, then detail the features that were extracted from the dataset, including user-generated features, system-related features, and context-related features. Finally, the clustering techniques are explained in detail as well as how they connect to the research questions outlined above.

Data Collection

In this work we address some of the gaps in the extant body of research by presenting a series of analyses from a large corpus of automatically logged phone-usage data, gathered by logging 3,043 individuals for a maximum duration of 3 weeks using the logging tool MobileDNA. MobileDNA records which apps are used, if multiple apps are used within a single session of usage, when notifications are received, if app use is initiated by a notification, and how much time a user spends interacting with each app. A session starts with unlocking the phone and ends when the screen is locked again. Only apps that actively engage the screen are logged, for instance playing music with an inactive screen is not considered to be time a user spends in an app.

User-generated features

App distribution. This set of features quantifies the distribution of app usage for a particular user across the apps with the highest frequency among all users. The 15 highest frequency apps are identified and each feature quantifies the proportion of times this app is used by a particular user relative to the total number of app usages (see Figure 1). An additional feature that captures the proportion of all other apps is included. The distribution of app usage provide a fundamental representation of the app repertoire of an individual user. Furthermore, comparing app repertoire across the same set of high frequency apps allows for the identification of specific niches of app usage that are idiosyncratic to a particular subset of users (Dimmick, 2002; Lin, Zhang, Jung & Kim, 2013).

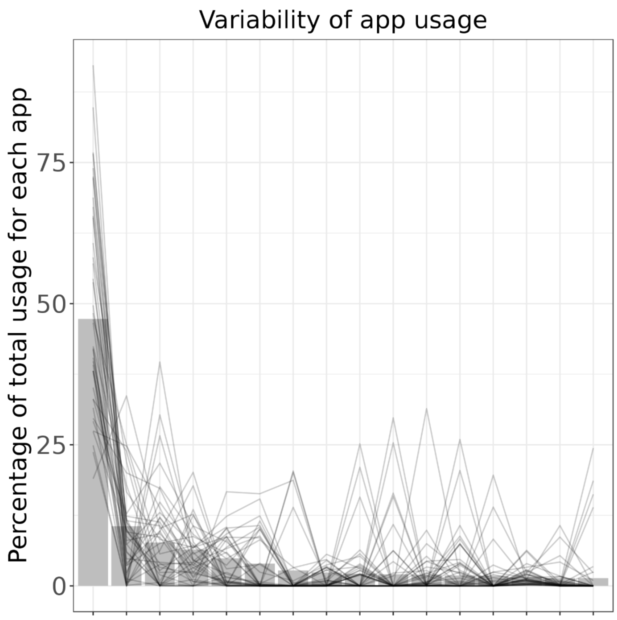

Figure 1: Variability of app usage across the 15 most frequently used apps. The bars show the mean app usage across all users, while the lines show the individual distributions for a small sample of users. The left-most bar indicates the proportion of all apps not in the top 15.

Multi app trains. Series of at least two apps that were used within a single session of usage were identified as trains of apps and the 150 most frequent trains across all users were extracted. A set of 150 features were built for each user, where each feature coded the proportion of that user's sessions that contained each high-frequency train. High-frequency trains of apps extend and enrich the concept of app repertoires to include sequences of multiple apps that frequently occur in order. Additionally, a high proportion of identified trains is a hallmark of habituated behavior where the progression from one app to the next within a session of use is nearly automatic.

Multi app sessions. The multiple applications feature identifies sessions that contain three or more app events. This is represented as the proportion of those sessions out of total sessions per user (left panel of Figure 3). Sessions containing many applications are the inverse of highly fragmented usage behavior and thus provides a clear measure of the degree of fragmented usage.

Repeated app sessions. The repeated app sessions feature identifies sessions that contain more than one instance of a single app with at least one app interspersed. This is represented as the proportion of a user's sessions that include such a repeated app (middle panel of Figure 3). Repeated app usage within a session also is indicative of habituated behaviors, specifically looping, repetition, and possibly perseveration on a single app.

Whatsapp repeat sessions. The Whatsapp repeat sessions feature is the proportion of Whatsapp sessions where a user engaged with the app two or more times in the session (right panel of Figure 3). This feature is indicative of conversational patterns of smartphone use.

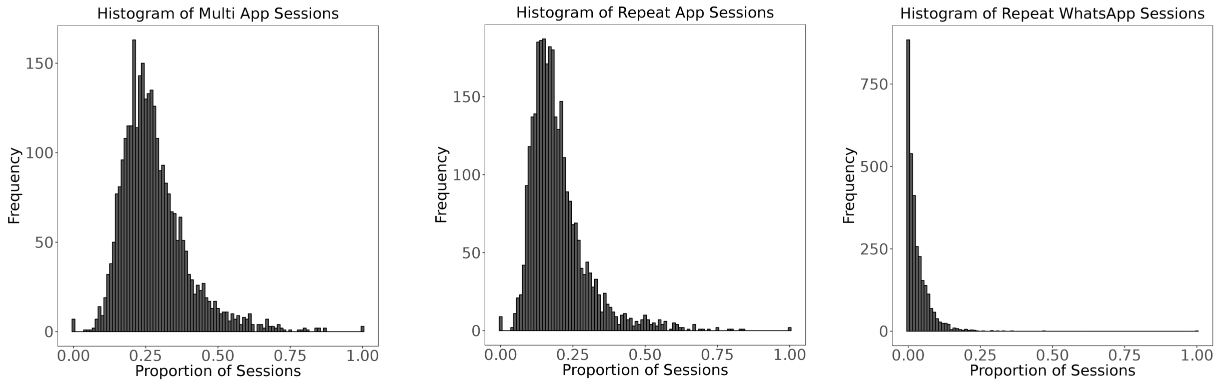

Figure 2: Histograms of average proportion of sessions that contain multiple apps (left), repeat an app at least ones (middle), and repeat WhatsApp within a session.

Between session duration. This feature quantifies the average time between sessions for each user. With the very long breaks between usage during sleep and other activities, the median provides a more clear measure of central tendency for such highly skewed distributions than the mean. As with multi-app sessions, between session duration is inversely related to fragmented usage: very low durations between sessions is indicative of high fragmented usage patterns (cf. Oulasvirta, Rattenbury, Ma & Raita, 2012).

Relative frequency of usage. The frequency of app usage is measured as the ratio between the total number of times using any specific app relative to the total amount of time app usage recorded. A high frequency of app usage per hour is likely to be an indicator of a highly fragmented usage pattern.

Relative duration of usage. The duration of app usage is measured as the ratio between the total amount of time spent using any app relative to the total amount of time that app usage is recorded. A low average duration of app usage per hour is likely to be an indicator of a highly fragmented usage pattern.

The relative frequency and duration of usage, as well as the between session duration are expected to be highly related but in potentially complex ways. People who initiate more app events are likely to spend more time using apps and have lower amount of time between sessions. However, longer sessions of using a single app are likely to result in lower frequency of app use. The relationships between these three measures is visualized in Figure 3.

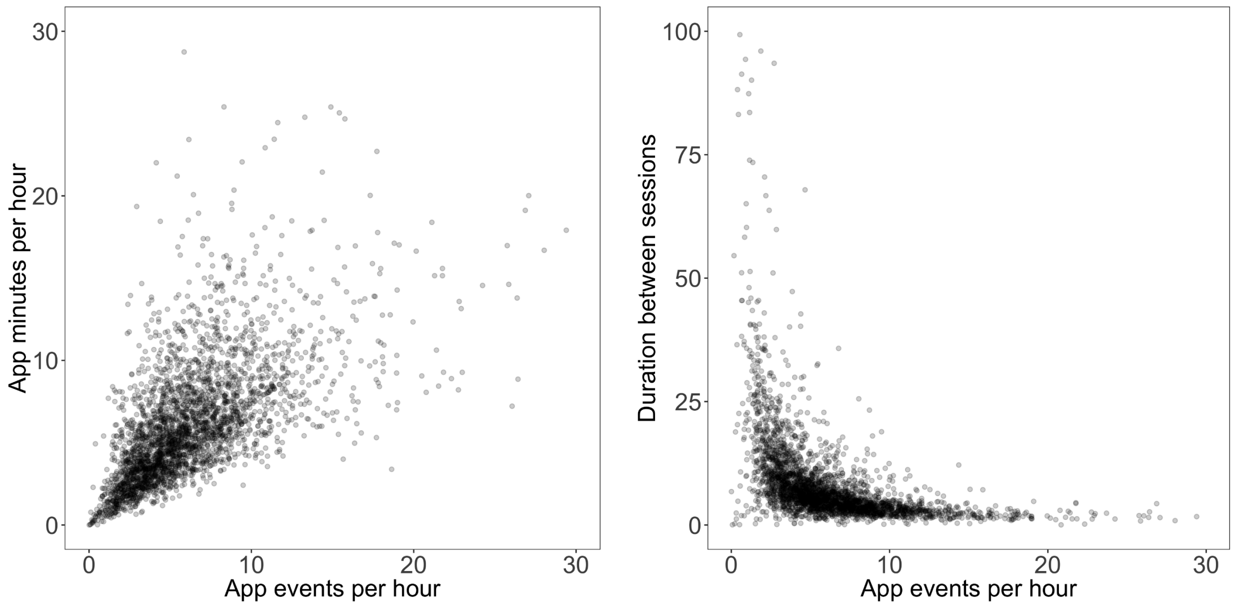

Figure 3: The relationship between the mean number of apps used per hour, the mean number of minutes per hour, and the median duration between sessions. The number of apps and minutes per hour are positively related (left panel) while the duration between sessions has a negative relationship and is highly non-linear (right panel).

System-related Features

Notifications. All app usage can be initiated either by the user without any prompt from the system, or by an interaction starting by engaging with a notification that alerts the user. For each app that is used, we can evaluate if the session in which it was used was initiated by a recent notification or by the user without prompting by the system. The initiating source of an interaction could have a significant impact on the types of usage patterns that are exhibited. For example, when splitting the user-generated features by the source of the interaction, the set of apps more likely to produce notifications that trigger usage are likely to be overrepresented for some users in the set of highest frequency apps, producing a dependency between the source of an interaction and the resulting app repertoire for that user. The same split based on notification may also be reflected in the patterns of habituation and fragmentation of use. The effect of splitting analysis by the source of the interaction, notification or intrinsic initiation, is evaluated in the second analysis.

Context-related Features

First and last contexts. The first and last interactions within a day are likely to be relatively unique usage patterns given the previous or expected long break from useage. They are likely to be more habituated than standard usage, and potentially contain unique fragmentation patterns and app repertoires that might deviate from the normal usage pattern. We identify these first and last contexts as the first and last 30 minutes a user is actively using their smartphone. Due to the high frequency of phone use after midnight, we consider the 'beginning' of each new day to be 5:00 am. A critical question, as it was for the system-generated features, then emerges: does taking these unique times into account by evaluating the features independently in each context matter for understanding smartphone usage?

Analyses

In order to address our three research questions, we compare the stability of clustering solutions across different sets of features. In all cases, the 3,043 users are clustered into groups based on their similarity across multiple features. The clustering solutions were derived by using affinity propagation (Frey & Dueck, 2007), where the number of clusters is not predefined nor are clusters biased to be equal size, but users that exemplify sets of users are iteratively found and all users are clustered based on their similarity to those exemplar users. Before the clustering solution is found, the dimensionality of the features defining users is reduced to 20 dimensions using principal components analysis (Pearson, 1901). This has the dual benefit of eliminating correlations between features as well as normalizing the feature dimensions.

The impact of changing the set of features is evaluated by measuring how much the clustering solution is different between the two sets of features. Adding features that are redundant with other features will not change the similarity between users that drives clustering, but adding or removing features that provide previously unavailable information that changes the similarity between users should dramatically impact clustering similarity. The similarity is quantified using the adjusted Rand index (Hubert & Arabie, 1985), which measures similarity of clustering solutions and varies from -1 to 1, with 1 indicating perfectly matching solutions and 0 corresponding to near chance similarity.

With respect to the person-, system-, and context-specificity of patterns, the first question concerning the stability of the conceptually relevant features within one individual, we compare the clustering solutions where the features are extracted separately from the first half and second half of each user's data. The question focusing on system-specificity is answered by looking at stability when accounting for system-related features such as notifications. To answer this question, we compare a clustering solution based on features from the full set of data to a clustering solution that includes two sets of features, including the features extracted from notification-initiated and the user-initiated usage separately. Finally, the question concerning context-specificity concerns stability in patterns when accounting for context-related features. For this analysis, we compare the clustering solution based on features from the full set of data to a clustering solution that incorporates three sets of features: features from the first 30 minutes of usage each day, from the last 30 minutes of usage each day, and from all other usage. In all analyses, the full set of features are extracted for each subset of the data.

Results

User-generated features

App distribution. The overall pattern of app frequency (Figure 1) matches the findings of Jung, et al. (2014) that the most used apps are communication and social media apps. The most frequently used application was Whatsapp (10.53%), followed by Facebook (7.69%), Chrome browser (6.46%), and Facebook Messenger (4.55%). Furthermore, the proportion of an individual app use varies drastically across individuals. For each app in the most frequent 15 there are users who either did not use this app at at all and some users for whom the app accounts for more than 30% of their total app events. Interestingly, nearly half of the app events are not one of the 15 most frequent apps, suggesting a long-tail of low frequency apps that are potentially more unique to individual users.

App sequences and trains. The majority of sessions contained only one distinct app in the session (55%). This, more than any other measure, highlights the prevalence of fragmented user behavior. Of the remaining sessions with two or more distinct apps, 54% used only two distinct apps in a session. However, unlike the distribution of app events, the distribution of the trains with two or more apps is much more uniform with a low peak at which the most frequent train only accounts for 2% of all trains.

This pattern of results presents a paradox. Most users are using their phone in short sessions and the significant number of users engaging with these short, frequent trains supports the idea that there is some similarity of behavior among basic app engagement. For example, the second most frequent multiple-app train was using Whatsapp then the Facebook app, and 38% of users displayed this behavior at least once. However, this pattern falls off sharply and the low frequency of each individual train across all trains suggests idiosyncratic patterns of use that might not be best evaluated at the level of individual apps.

User-generated features	**Median**	**Mean**	**SD**
Proportion of sessions with multiple apps	0.26	0.28	0.11
Proportion of sessions with repeated apps	0.18	0.20	0.10
Proportion of sessions with multiple Whatsapp	0.017	0.031	0.043
Median minutes between sessions	5.58	9.26	22.11
App events per hour	5.45	6.51	9.37
Minutes of use per hour	5.48	6.23	4.32

Table 1: Descriptive statistics of the user-generated features that summarize broader patterns of usage.

Other user-generated features. Table 1 summarizes the measures of central tendency and variance for the remaining user-generated features. We see reasonably high proportions of sessions containing multiple apps (26%) or repeated apps (18%). Though the distribution of these measures do not seem particularly skewed, the relatively high standard deviation indicates a high degree of variability across individuals. In contrast, the proportion of sessions with multiple instances of the Whatsapp app, indicative of a conversational mode of smartphone usage, was startlingly low at 1.7% with a variance that had very few users with more than 10% of such conversational sessions based around Whatsapp. Finally, the measures of usage: time between sessions, app events per hour, and minutes of use per hour were relatively highly skewed and highly variable across individuals. Figure 3 illustrates the relationship between these features and shows that generally the number of app events and the number of minutes of use per hour are positively related while the relationship between time between sessions and app events per hour is a highly non-linear negative relationship.

Clustering Results

The similarity between clustering solutions is shown in Table 2. The overall results suggest the stability of clustering solutions across different sets of features is low. This is most apparent for the similarity between the clustering solutions based off the features extracted from the first and second half of each user's data (similarity of 0.018), which is very close to the adjusted Rand index for random similarity (similarity of 0). However, the similarity between the features from the full dataset and the notification-split and time-of-day-split datasets are marginally higher.

First feature set	**Second feature set**	**adjusted Rand index**
Features from first half of dataset	Features from second half of dataset	0.018
Features from full dataset	Two sets of features from usage initiated by notifications or not	0.23
Features from full dataset	Three sets of features from datasets split by time of day	0.13

Table 2: The similarity scores between clustering solutions based on five different sets of user-generated features.

Overall, the similarities between the clustering solutions are closer to chance than perfect matches, however a clear trend emerges: the similarity between the full data and notification-split data is more similar than the full data and time-of-day-split data, which is more similar than the similarity between the first half and second half of a users data.

Our results indicate support for many existing hypothesized usage patterns but present a more complete and nuanced view of how people use smartphones, and identify relevant user-level, system-level and context-level variability in patterns of use. The analytical procedures used in this article provide hands-on instructions for media effects scholars to include the identified patterns in their own analyses.

Discussion

The descriptive analysis of the user-generated features shows strong support for the theoretical constructs of an individual app repertoire and generally high fragmentation of usage patterns, as well as moderate evidence of habituation behaviors. The importance of individual app repertoire is most clear from the variability in the distribution of app usage across the most frequent apps, though there are indications of general usage patterns that are shared by most individuals. Fragmentation of usage is highlighted in the high proportion of short trains of consecutive apps within a single usage session as well as the measures of average usage time and average number of app events per hour. However, these measures do indicate significant variability in fragmentation across users. The evidence for habituation is less clear. Relatively high measures of repeated app usage within a single usage session suggests habitual app-checking behavior and these repeated sessions account for most of the sessions containing more than one app. However, the pattern of many low frequency trains of apps suggests higher variability in usage patterns that are not ritualized or habitual.

The clustering results indicate a more nuanced pattern. The first cluster result, that splitting a user's data into two sets, produces clustering solutions that are negligibly more similar than random chance, is potentially troubling for the the hypothesis that these user-generated features provide meaningful, stable constructs for evaluating smartphone use. One clear limitation with this analysis, however, is that after splitting each user's data in half, only 9 or 10 days were included in each half per user. It is possible and even likely that this is not enough data to derive stable measures of usage behavior. This viewpoint is strengthened by looking at the similarities between the features from the full dataset and the notification-split and time-of-day-split datasets. These clustering solutions, which are all based on the full set of 21 days from each user, show higher similarity values, though far from a perfect match.

Furthermore, the lower similarity to the basic clustering solution for the time-of-day-based features than the notification-based features suggests the time-of-day features capture variability that is not present in the overall features. One distinct possibility is that this additional variability is noise due to having smaller amounts of data in the first and last periods of each day. However, another possibility is that the first and last times a person uses their phone might be markedly different than how they use it throughout the day and this provides critical information about how to define similarity across users. This view is consistent with the prediction that patterns of habituated usage (Bayer, Campbell, & Ling, 2016; Oulasvirta et al., 2011), particularly during times when habituation would be more likely, are critical to understanding smartphone usage.

Overall, we find evidence of user-generated features of smartphone use that provides insight on the degree of user habituation and fragmentation as well as the unique app repertoire of individuals. The clustering analyses suggest that the degree to which these features can be used to identify stable clusters of users is influenced by the degree to which behavior is driven by system-regulated notifications or during habituated early and late phone usage. However, clustering solutions based on these user-generated features may be unstable, at least when extracted from relatively short durations. Further analyses based on larger durations of data collection are necessary to address this question and can shed further light on the specific differences of user-generated features extracted from specific system-related and context-related events. Finally, these features present a clear opportunity to utilize these features to better understand smartphone usage and predict aspects of well-being and mental health.

References

Bayer, J. B., & Campbell, S. W. (2012). Texting while driving on automatic: Considering the frequency-independent side of habit. *Computers in Human Behavior, 28*(6), 2083-2090.

Bayer, J. B., Campbell, S. W., & Ling, R. (2016). Connection cues: Activating the norms and habits of social connectedness. *Communication Theory, 26*(2), 128-149.

Collier, R. (2016). Mental health in the smartphone era. CMAJ : Canadian Medical Association journal = journal de l'Association medicale canadienne, 188(16), 1141-1142.

David, P., Kim, J.-H., Brickman, J. S., Ran, W., & Curtis, C. M. (2015). Mobile phone distraction while studying. *New Media & Society, 17*(10), 1661-1679.

Deng, T., Kanthawala, S., Meng, J., Peng, W., Kononova, A., Hao, Q., & David, P. (2018). Measuring smartphone usage and task switching with log tracking and self-reports. *Mobile Media & Communication*, 2050157918761491.

Dimmick, J. W. (2002). *Media competition and coexistence: The theory of the niche*. Routledge.

Dimmick, J., Kline, S., & Stafford, L. (2016). The Gratification Niches of Personal E-mail and the Telephone. *Communication Research, 27*(2), 227-248. doi:10.1177/009365000027002005

Frey, B. J., & Dueck, D. (2007). Clustering by passing messages between data points. Science, 315(5814), 972-976.

Gouin-Vallerand, C., & Mezghani, N. (2014, September). An analysis of the transitions between mobile application usages based on markov chains. In

Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (pp. 373-378). ACM.

Hao, Y., Wang, Z., & Xu, X. (2016, June). Global and personal app networks: characterizing social relations among mobile apps. In *Services Computing (SCC), 2016 IEEE International Conference on* (pp. 227-234). IEEE.

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of classification*, 2(1), 193-218.

Hsiao, K.-L., Shu, Y., & Huang, T.-C. (2017). Exploring the effect of compulsive social app usage on technostress and academic performance: Perspectives from personality traits. *Telematics and Informatics, 34*(2), 679-690.

ITU. (2017). *ICT Facts and Figures 2017*. Retrieved from https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2017.pdf

Jenaro, C., Flores, N., Gómez-Vela, M., González-Gil, F., & Caballo, C. (2009). Problematic internet and cell-phone use: Psychological, behavioral, and health correlates. *Addiction Research & Theory, 15*(3), 309-320. doi:10.1080/16066350701350247

Jung, J., Kim, Y., & Chan-Olmsted, S. (2014). Measuring usage concentration of smartphone applications: Selective repertoire in a marketplace of choices. *Mobile Media & Communication, 2*(3), 352-368.

Lanaj, K., Johnson, R. E., & Barnes, C. M. (2014). Beginning the workday yet already depleted? Consequences of late-night smartphone use and sleep. *Organizational Behavior and Human Decision Processes*, 124(1), 11-23.

Licoppe, C. (2004). ‘Connected’ Presence: The Emergence of a New Repertoire for Managing Social Relationships in a Changing Communication Technoscape. Environment and Planning D: Society and Space, 22(1), 135–156. https://doi.org/10.1068/d323t

Lin, W. Y., Zhang, X., Jung, J. Y., & Kim, Y. C. (2013). From the wired to wireless generation? Investigating teens' Internet use through the mobile phone. *Telecommunications Policy, 37*(8), 651-661.

Ling, R. and Lai, C. (2016), Microcoordination 2.0: Social Coordination in the Age of Smartphones and Messaging Apps. *Journal of Communication, 66:* 834-856. doi:10.1111/jcom.12251

Oulasvirta, A., Rattenbury, T., Ma, L., & Raita, E. (2011). Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing, 16*(1), 105-114. doi:10.1007/s00779-0

Pearson, K. (1901). LIII. On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2(11), 559-572.11-0412-2. doi:10.1080/14786440109462720

Poushter, J., & Stewart, R. (2016). *Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies*. Retrieved from http://www.pewresearch.org/wp-content/uploads/sites/2/2016/02/pew_research_center_global_technology_report_final_february_22__2016.pdf

Ramirez, A., Dimmick, J., Feaster, J., & Lin, S.-F. (2008). Revisiting Interpersonal Media Competition. *Communication Research, 35*(4), 529-547. doi:10.1177/0093650208315979

Rettie, R. (2009). SMS: Exploiting the interactional characteristics of near-synchrony. *Information, Communication & Society*, *12*(8), 1131-1148.

Schnauber-Stockmann, A., Meier, A., & Reinecke, L. (2018). Procrastination out of habit? The role of impulsive versus reflective media selection in procrastinatory media use. *Media Psychology*, 1-29.

Spira, J. B., & Feintuch, J. B. (2005). The cost of not paying attention: How interruptions impact knowledge worker productivity. *Report from Basex*.

Twenge, J. M., Joiner, T. E., Rogers, M. L., & Martin, G. N. (2018). Increases in depressive symptoms, suicide-related outcomes, and suicide rates among US adolescents after 2010 and links to increased new media screen time. *Clinical Psychological Science, 6*(1), 3-17.

Vanden Abeele, M. M. P., Antheunis, M. L., & Schouten, A. P. (2016). The effect of mobile messaging during a conversation on impression formation and interaction quality. *Computers in Human Behavior, 62*, 562-569.

Vanden Abeele, M. M. P., Beullens, K., & Roe, K. (2013). Measuring mobile phone use: Gender, age and real usage level in relation to the accuracy and validity of self-reported mobile phone use. *Mobile Media & Communication, 1*(2), 213-236.

Vanden Abeele, M. M. P., De Wolf, R., & Ling, R. (2018). Mobile media and social space: How anytime, anyplace connectivity structures everyday life. *Media and Communication, 6*(2), 5-14.

Vanden Abeele, M. M. P., Schouten, A. P., & Verbrugge, K. (2016). Slacking by checking? A study of employees' perceived internet checking habit in relation to their perceived work efficiency.

Wajcman, J., Bittman, M., & Brown, J. E. (2008). Families without borders: Mobile phones, connectedness and work-home divisions. *Sociology, 42*(4), 635-652.

Williams, J. (2018). *Stand out of our Light: Freedom and Resistance in the Attention Economy*: Cambridge University Press.

Wu, D., Liang, S., & Tang, Y. (2017). Towards better understanding of app transitions in mobile search. *iConference 2017 Proceedings.* Retrieved from https://www.ideals.illinois.edu/handle/2142/96752