UNIVERSITY OF
BATH

University of Bath

Revisitation analysis of smartphone app use

Simon L. Jones
University of Bath
Bath, UK
s.l.jones@bath.ac.uk

Denzil Ferreira, Simo Hosio, Jorge Goncalves, Vassilis Kostakos
University of Oulu
Oulu, Finland
{denzil.ferreira, simo.hosio, jgoncalv, vassilis}@ee.oulu.fi

ABSTRACT
We present a revisitation analysis of smartphone use to investigate the question: do smartphones induce usage habits? We analysed three months of application launch logs from 165 users in naturalistic settings. Our analysis reveals distinct clusters of applications and users which share similar revisitation patterns. However, we show that much of smartphone usage on a macro-level is very similar to web browsing on desktops, and thus argue that smartphone usage is driven by innate service needs rather than technology characteristics. On the other hand, on a micro-level we identify unique characteristics in smartphone usage, and we present a rudimentary model that accounts for 92% in the variability of our smartphone use.

Author Keywords
Revisitation, smartphone use, habits, user behaviour

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
We present a large-scale longitudinal analysis that seeks to quantify smartphone application use habits. A growing number of analyses in recent years have sought to investigate and model how individuals use their mobile phones. This has ranged from studying how individuals charge and manage battery [12], download, install and use different applications [11], and how use varies with context [10]. In terms of *habits* (i.e. patterns, routines) in smartphone application use, previous work has sought to model and predict which applications people are likely to install [37], use [31], and for how long [13].

While previous work has repeatedly shown that individuals exhibit well-structured and predictable patterns in their smartphone use, we are still far from developing a theoretical understanding of why these behaviours arise.

This is important because it could provide guidance on the development of new mobile technology and interaction techniques, and on evaluating and comparing existing ones. For instance, while it is commonly accepted that mobile phones are "habit-forming" [25], it is not clear whether these habits are driven by the users themselves, their phones' affordances, or some combination of both. At the same time, little work has quantitatively compared how our smartphone use habits relate to our use habits for other technologies, particularly due to our lack of theoretical models of use.

We address both of these challenges simultaneously, by re-using an established methodology that has been extensively employed in prior literature, albeit in the context of web browsing: *revisitation analysis* [33]. Here we study how users revisit: (1) apps on their phone (i.e. how often do they return to a particular app?). This helps us quantitatively describe every individual app and the type of revisitation it attracts across users; (2) the phone itself (i.e. how often do they use apps on their phone?). This helps us to quantify individual users in terms of their overall revisitation habits; (3) applications in-session (i.e. how often they switch back and forth between active applications between unlocking and locking their phone?). This helps us to quantify backtracking and multitasking on a micro-level.

By re-using an established method to analyse revisitation patterns, we are able to compare revisitation behaviours that individuals exhibit on smartphones against those in desktop-based web-browsing. Our premise is that commonalities found across such different technologies are likely to reflect innate user habits, and therefore we are able to tease apart those habits that appear to be driven solely by smartphone technology.

Ultimately, we seek to make two contributions to the corpus on understanding mobile use. First, we attempt to investigate whether mobile phones are a "special" type of technology in the sense that they are habit-forming and shape our daily routines, or whether humans just exhibit the same habits of use across different technologies. The revisitation analysis method we employ focuses on detecting temporal patterns, hence in this paper we identify and compare habits that are temporal in nature (i.e. behaviours that are triggered after a certain time interval since the phone and its apps were last used).

Second, we provide the tools and techniques for conducting a systematic analysis of revisitation behaviour patterns, and in our case apply it to smartphone use. An ongoing discussion and self-criticism within HCI and Ubicomp literature is that studies tend to avoid re-use of research methodologies [19,20], thus making it hard to reliably compare results with previously published work, and to consistently build theory. In other words, by developing tools that enable us to consistently build on previous literature we facilitate the articulation of testable theories of technology use, and in our case of smartphone revisitation habits.

RELATED WORK

Temporal patterns in smartphone use

Previous work has extensively focused on how individuals use their smartphones and their applications. For instance, Yan *et al.* [36] found that the majority of mobile device usage is brief: 50% of mobile phone engagement (the time period between the user unlocking and relocking the device) lasts less than 30 seconds. More recent work has shown similar results [13], further focusing on micro-usage habits that entail using apps in short bursts of less than 15 seconds, suggesting a strong temporal nature in smartphone use. Similarly, Falaki *et al.*'s [11] study on diversity in smartphone battery usage reported finding short-term application usage (10-250 seconds) among their participants, and Böhmer *et al.*'s [3] large-scale study on mobile application usage revealed that mobile phone owners use their device for an average of 59 minutes daily, with the average application session lasting 72 seconds. Oulasvirta *et al.* [26] demonstrated how attentional limitations emerge in the mobile contexts and result in short bursts of smartphone interaction.

Considering routine, and focusing on overall mobile phone users' habits, Oulasvirta *et al.* [25] suggest that mobile phones are "*habit-forming*" devices, highlighting the "checking habit: brief, repetitive inspection of dynamic content quickly accessible on the device." This habit was found to comprise a large part of mobile phone usage, and follow-up work [13] has argued that the checking habit is one of the behavioral characteristics that leads to mobile application micro-usage, which is subsequently manifested as short bursts of interaction with applications. However, it remains unclear whether other temporal patterns of behaviour exist, and whether they can be attributed to the technology, to humans, or a combination of both.

For instance, rapid switching between applications is not exclusively mobile: multitasking and the management of multiple tasks occur on desktop computers as well [8,15]. These studies found that desktop users were focused on goals but within these goals they embedded a set of secondary goals that were completed through multitasking, albeit not without impacting the users' attention and workflow [21]. This observation led to the technical advances within multitasking and the options that almost all multipurpose technical devices support: being able to seamlessly switch from one application (and often from one task) to another.

Further elaborating on cross-technology habits, previous research [2] has highlighted how smartphone users would 'control' their own use of communication by checking email and other communications on their smartphone, and then waiting to reply until they could access a larger device such as their laptop. This resonates well with findings reported in literature regarding brief bursts of micro-usage interaction: mobile devices seem to be often used for brief checking.

Modelling and prediction

Modelling mobile phone application use enables prediction of which applications will be launched next, or which likely actions will be taken next by the user. This could help in e.g. automatically promoting certain application shortcuts to the phone's home screen [31], or optimising the phone's internal memory management [36]. Several approaches have been trialed to create such application use models. For instance, Huang *et al.* [16] considers contextual information about last used application, time, and location and user profile to predict the application that will be used next. The results showed that a regression model works best by incorporating identified sequences of application use in predicting the next application. This suggests a strong sequential nature in application usage on smartphones.

Leroux *et al.* developed a mobile framework that creates *context profiles* by monitoring application use, day of week, and the user's speed and location [18]. Although using artificially created data, the resulting profiles match real life situations, such as "at work" or "commuting," and can be used to infer a set of applications the user is likely to use. In a similar study, a much wider set of context attributes, such as location, cellular activity (SMS, calls), personal schedule, battery status, screen and status, was used in addition to application use to provide a naturalistic and reliable model for making the predictions to create a dynamic home screen to reduce the time to find an application [31].

In summary, we found that an emerging theme in predicting applications is analysing the sequential and temporal nature of application use, and especially the suitability of using previous application launch history as the predictor of next applications to be launched. The analysis may focus on a micro-scale, i.e. which applications were used right before the predicted applications [16], consider longer periods of time during which an application is "trending" for the user [36], or establish clusters of applications that get used together often [36]. Our work aims to complement the existing approaches by analyzing temporal and sequential application usage in depth. Next we summarise an established methodology that can enable us to study the temporal nature of application usage on smartphones at varying timescales. It also provides flexibility in terms of

making either users or applications the focus of the analysis.

Revisitation analysis
Previous research has shown that web page revisitation is one of the most frequent actions in computer use [1]. As a consequence, many researchers have sought to gain a deeper understanding of web revisitation behaviour and improve mechanisms for web navigation such as back buttons [24] and history mechanisms [32]. A significant number of studies in this area have collected logs of web activity to analyse browsing patterns.

A study on web revisitation by Tauscher & Greenberg [32] found a mean revisitation rate of 58%, i.e. on average, 58% of the web pages visited by a user had already been visited by that user at some point in the past. A subsequent study by Cockburn & McKenzie [7] reported web pages revisitation rates of 81%. More recent studies have shown that tabbed browsers – which provide the ability to visit multiple pages in parallel – have further altered browsing behaviour [38].

Obendorf *et al.* [24] developed a theoretical basis for revisitation analysis by distinguishing between different types of revisitation based on their inter-visit temporal duration: *Short-term*: these are instances when a user revisits a particular web page within a short period of time. This behaviour is akin to backtracking or undo. *Medium-term*: this behaviour reflects an intention to re-utilize or observe a particular page. *Long-term*: this behaviour is described as rediscovering or reusing a particular page. Therefore, Obendorf *et al.* argue that revisitation patterns are closely tied to purpose and intention. More relevant to our own work, Adar *et al.* [1] introduced the concept of the 'revisitation curve.' These are normalized histograms of inter-visit times, consisting of 15 exponentially-spaced bins, that can be used to characterise the revisitation pattern for a particular web page. Thus, a revisitation curve is effectively a smoothed histogram that provides a finer resolution than Obendorf *et al.*'s triadic categorisation.

Using revisitation curves it has been demonstrated that different categories of web sites invite different types of revisitation behaviour across all users [1]. For instance, search engines are revisited at extremely short intervals, webmail pages at medium intervals, and entertainment or hobby web pages at hybrid (combinations of short and long) intervals. Thus, revisitation curves are an ideal empirical lens for studying habits, as they are mathematically defined, empirically derived, and reliably comparable due to their quantitative nature.

Using revisitation curves has also made it possible to analyse the behaviour of individual users. Pushnyakov & Gusev [29] demonstrated that, much like web pages, web users have distinct revisitation curves which describe their web browsing behaviour. When considering web browsing on smartphones, Tossell *et al.* [33] showed that smartphone browsing sessions are three times shorter than desktop browsing sessions in terms of duration and pages visited, while web revisitation rates on smartphones are lower than desktops. They conclude that browsers on smartphones have largely given way to 'Native Internet Applications', suggesting that users increasingly access the Internet via native applications while using browsers for ad-hoc searching and medium-term revisitations. Finally, as in our own work, they propose that users' revisitation curves can be used to capture distinct user profiles and peculiarities.

The strong similarities between web browsing on desktops and native application use on smartphones [33] and the established methodology for analysis of web browsing behaviour [1], has lead us to consider revisitation analysis as a valuable approach for studying the use of smartphone applications. We hypothesize that different applications will have different revisitation patterns based on their purpose and content, and that users can be characterised based on their revisitation curves.

METHOD
We conducted a study using the Google Play application store as a platform for experimental data collection. We built a free application, Securacy, that allowed participants to monitor overall network activity on their mobile phones [14]. Securacy provides insight into the network connections that are established behind the scenes when apps are launched, including the use or otherwise of secure network connections and the geographical destination of the app data. Participants could choose to share the collected data with us, including the logs of which applications were launched and at what time. Participants were fully informed about the authors of the application and the scientific purpose of the experiment, and gave explicit consent when activating the software.

After deploying our software on the app store, we publicly promoted and advertised it on Facebook and Twitter in an attempt to reach a diverse population sample. We offered no additional compensation beyond the free use of the app, which offered users a useful function in itself. 165 participants downloaded the Securacy app and provided data using their Android smartphones in naturalistic settings, over a period of 3 months. Tablet devices were excluded from the analysis presented in this paper. 41% of participants were from European countries, 53% North America, and 6% others (e.g., Australia, South-America).

For our purposes, the application used the AWARE[1] framework to log a new record every time a user launched an application. The entry included a timestamp (in the time zone of the participant), a unique ID for the participant, a unique package name assigned to the application by the developer, and the localized application name.
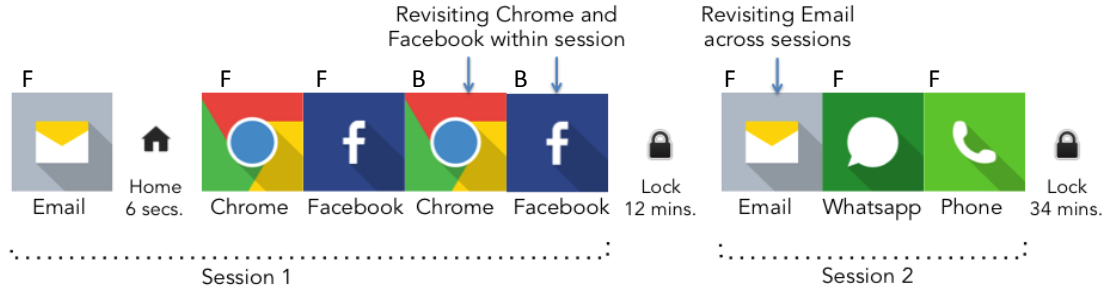
---

[1] www.awareframework.com

Figure 1. Example chain of application use across different sessions. An app may be visited for the first time in a session, tagged as a forward action (F), or revisited within a session (e.g. Chrome and Facebook), tagged as a backtracking action (B).

Our Securacy application interrupted use of other apps only under two conditions: when the user installed an application for the first time, and when the user removed an application. This meant that the application's functionality did not affect normal use of the phone, and that the app for data collection appeared infrequently in our app traces.

During the study we captured 199,859 application launch events for 1,527 unique applications across all users. In our analysis we ignore apps that are core to Android OS (e.g. the home screen and system UI), or a means to launch other applications. We found that many of the participants had installed skinned UIs and application launchers to replace parts of the standard Android OS. For these users we also ignore such system-wide applications.

Data Filtering
Since we are focusing on revisitation patterns, and to avoid outliers' bias in our data, we removed the users with fewer than 10 days of logged activity from our analysis. We also only analyse the revisitation patterns for apps that are installed by more than a single participant within our dataset (approximately 1400 apps within our dataset). Figure 2a shows the total number of app launches for each of our participants. Similarly, we constructed Figure 2b to determine the popularity of applications. Within the study there were few apps that were opened many times (approximately 100 apps being used more than 100 times each), and many apps that were opened only a few times (approximately 1000 apps opened less than 10 times each). The top 10 most frequently used apps (excluding filtered apps) were: Whatsapp, Chrome, Phone, Facebook, Google Hangouts, SMS/MMS, Gmail, Contacts, Play Store and Desk Clock.
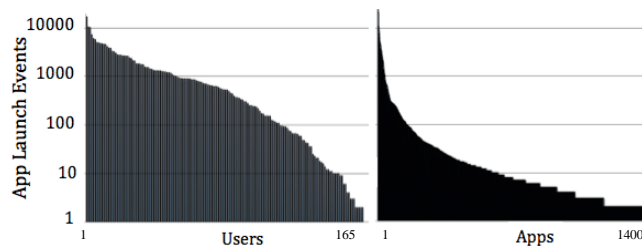

Figure 2a. Number of logged app launch events per user (left);
Figure 2b. Number of logged launches per app (right)

While users greatly varied in how much they use their phone (Figure 2a), and apps greatly varied in how popular they are (Figure 2b), both distributions seem to follow a pattern that suggests that the wide range in frequencies is not due to outliers but to a skewed distribution.

Deriving Revisitation Curves
The key component of our analysis is the revisitation curve [1] representing the number of times that an app is revisited within a predefined time interval. We construct a revisitation curve for a certain application by considering the duration between revisits to that application by users (Figure 1). Similarly, a revisitation curve for a particular user is constructed by considering the duration in-between launching *any* app on their phone.

Following the methodology in [1], we use an exponential scale for revisit interval bins. The following 15 bins were used: 1, 2, 4, 8, 16, and 32 minutes; 1, 2, 4, 8, 16, 32 (i.e. 1.3days), 64 (i.e. 2.6days), 128 (i.e. 5.3days), 256 hours and above (i.e. >10.6days). A revisitation curve characterizes an app by its 15-dimensional vector, where each dimension corresponds to the frequency of revisits within the corresponding bin. These curves are like a "signature" of users' behaviour in launching a given smartphone application.

We found that our dataset contained instances of users being interrupted by 'pop-up' applications or processes within apps (e.g. invoking keyboards, cameras, barcode scanners etc.). To mitigate the effects of these pop-ups in our analysis, any return to an app from a recognized pop-up app within 30 seconds is not treated as a revisit when constructing the revisitation curve. Pop-ups that appear for longer than 30 seconds *are* treated as significant switches away from an application, given that 50% of smartphone engagement itself lasts less than 30 seconds [36].

Figure 3 illustrates the revisitation curve for the popular messaging application WhatsApp. The curve indicates the frequency or probability (y-axis) of observing a revisitation to WhatsApp within a particular time period (x-axis). In this case the data is aggregated across all users, since all devices share the same package name per application. For instance, we observe that there is a 16% chance that once an average user launches WhatsApp, they will launch it again within 2-4 minutes (Figure 3). Similar curves were calculated for all applications in our dataset. Additionally, we are able to

calculate a revisitation curve for each user, as suggested in [29]. In this case the curve shows the probability the user, having just used an application, will use any other application within a particular time period.

For this analysis we use the same 15 temporal bins. Figure 3 also shows the revisitation curve for a single user from our study. In this example there is a 20% chance that the user will launch an application within 4-8mins following a previous application launch, and approximately 8% chance that they will wait 8-16hrs.
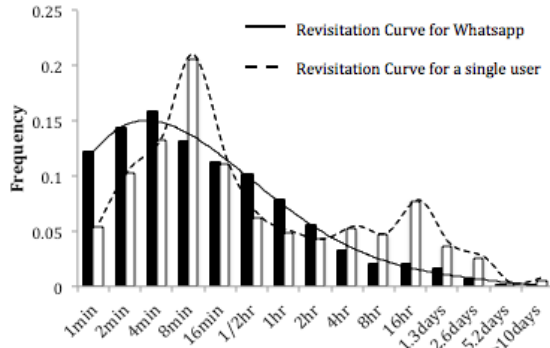

Figure 3. Revisitation histogram and curve for a single app (WhatsApp) and for a single user across all apps.

Clustering

The next step in our analysis consists of identifying clusters of similar applications or similar users. Each distinct application or user is described by a respective vector of 15 values corresponding to the 15 bins of their revisitation curve. By treating these values as features we are able to apply clustering algorithms to identify clusters with similar features. We used the k-means clustering algorithm to identify clusters of applications and clusters of users. Thus, we are able to identify applications and users that share similar revisitation curves. We iteratively apply k-means for a varying number of clusters and use within-groups sum of squares to build an elbow graph in order to identify the optimal number of clusters.

Revisitation analysis in-session

The final part of our analysis considers revisitation behaviour regarding applications that the user has recently launched since unlocking their phone. This analysis focuses on single "sessions" of use: for each session we only consider the set of applications that were used between unlocking and locking the phone. For every such session (performed by an individual user) we consider whether an application is launched for the first time (marked as "F" for forward), or whether it was launched previously within the session (marked as "B" for backtracking). For instance, consider a user who unlocks their phone and uses the following applications in sequence: Email, Chrome, Facebook, Chrome, Facebook (Figure 1), before re-locking their phone. This sequence of 5 application launches can be represented as: FFF-BB. Our encoding shows that the user used three distinct applications in succession (FFF), then used two applications that had already been used in the session (BB). For each string we can calculate its length and the ratio of B or F characters. Finally, we used regular expressions to group these strings into pre-determined categories. These categories emerged from qualitative analysis using two independent coders.
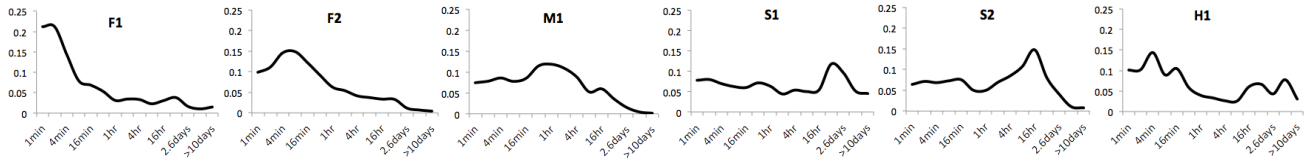
Cluster Label	Description	Centroid Revisitation Curve	Example Apps	Cluster Size (# Apps)
F1	Fast		Google Play Store, Facebook Messenger, InoReader, Chrome Beta, BlackBerry Messenger, Reddit, Okcupid	32 (13%)
F2	Fast		Chrome, Whatsapp, Facebook, Google Hangouts, SMS/MMS, Viber, Youtube, Contacts+, Google Maps, Firefox, Spotify, Skype, Snapchat, Xperia Conversations, Line, Reddit News, Telegram Messenger, Music, Falcon Pro	82 (33%)
M1	Medium		Phone, Gmail, Contacts, Email, Dialer, Clash of Clans, Instagram, Outlook, Yahoo Mail, Opera Browser	47 (19%)
S1	Slow		Gallery 3D, Calendar, Camera, Twitter, Calculator, Clean Master (Speed Booster) Runkeeper Pro, Flipboard, Google Play Services, Mobile Bank, Mobile Weather, Flickr, Google Doc Editor, Tumblr, Quick Office, Google Translate	30 (12%)
S2	Slow		Settings, Desk Clock, Organiser, Tinder, Plants vs. Zombies 2, Clash of Lords, Titanium Backup, Hot or Not, Control Panel, Candy Crush Saga, Castle Clash	40 (16%)
H1	Hybrid		Evernote, Google+, Google Docs, MusicBox, Adobe Reader, 9gag, Video Player, Meo Remote, Waze, Dictionary, Opera Mini	21 (8%)

Table 1. The 6 clusters of applications based on their revisitation curve. For each cluster we provide a label and description, a visualisation of its centroid revisitation curve, some representative apps from the cluster, and the popularity of the cluster

RESULTS

Revisitation Curves: Applications

Table 1 shows the 6 application clusters and the revisitation curve for each cluster centroid (i.e. the prototypical revisitation curve for that cluster). Since revisitation curves consist of 15 bins, we conceptually splits those bins into three 5-bin segments: fast, medium, and slow, as proposed in [24]. Each cluster is then characterised based on where its frequency modality lies (fast, medium, slow), or whether it has a bimodal distribution (hybrid).

We found that apps in cluster F1 had a peak revisitation at 1-2 minutes. The types of apps in F1 include *instant messaging* (Facebook Messenger, Blackberry Messenger), *app store*, *social media* and *dating*, and *web browser*. The F2 cluster had peak revisitation at 8-16 minutes and also included *instant messaging* (Whatsapp, SMS, Viber, Skype, Snapchat, Telegram Messenger, Google Hangouts, Line), *browsers* (Chrome, Firefox), and *social media* (Facebook, Reddit, Falcon Pro (Twitter client)).

On the other hand, applications in cluster M1 peaked their revisitations at 1-16 hours. The types of apps in this cluster include *email* (Android mail client, Gmail, Yahoo Mail, Outlook) and *phone communication* (Phone, Contacts, Dialer).

The applications with slow revisits were grouped into two clusters. In cluster S1 applications were revisited mostly at 2-3 days, and included *utilities* (e.g. Calculator, Camera, Speed Booster, Calendar, Mobile Banking, Weather, Translator, Document Editors), *multimedia* (Photo Gallery, Flickr, Magazines (Flipboard app)) and *health and fitness trackers*. Applications in cluster S2 peaked at 24 hour revisits, and included *games* (Castle Clash, Clash of Lords, Plants vs. Zombies, Candy Crush Saga), *dating* (Hot or Not, Tinder) and *phone settings* (Control Panel, Settings, Alarm Clock). Finally, in cluster H1 we found apps with peaks at 4 minutes, 16 minutes, 1 day, and 10 days. These are apps which that are typically used on a weekly or fortnightly basis, but the sessions in which they are used contain multiple shorter revisits. This cluster includes apps for *documents* and *notes* (Evernote, Google Docs, Adobe Reader), *GPS and Satnav* (Waze).

Revisitation Curves: Users

In our analysis we excluded 36 users because they did not have enough revisits (fewer than 10 app revisits in total) or did not participate in the study for longer than the maximum bin (10 days). For the remainder of users we
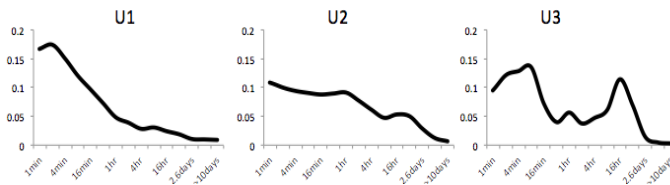

Figure 4. Centroid revisitation curves for the three clusters of users - Checkers (U1), Waiters (U2), and Responsives (U3)
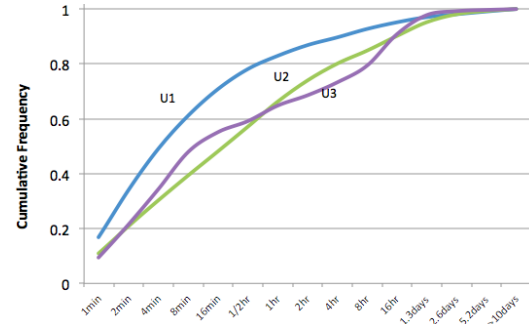

Figure 5. Cumulative probability distribution for each user type: checkers (U1), waiters (U2), responsives (U3).

identified a substantial trichotomy when clustering their user profiles (Figure 4).

Checkers (U1). These users exhibit brief revisit patterns heavily skewed towards fast revisitation (less than an hour). Checkers account for 44% of users.

Waiters (U2). These users exhibit longer revisit patterns uniformly distributed between short-medium revisitations (between 1min and 4hrs) and long revisitations (2hrs - 3 days). Waiters account for 46% of users.

Responsives (U3). These users exhibit sometimes brief and sometimes long revisit patterns. Responsives account for 10% of users.

Figure 5 shows the cumulative probability distribution for all user types. We observe that *Checkers* (U1) have a 60% probability of launching an application 4-8 minutes after launching a previous one – across the whole duration of the study. This ratio becomes 80% at the ½ hour mark. We also observe that the 1 hour and 8 hour marks are pivot points for *Waiters* (U2) and *Responsives* (U3): between those marks *Waiters* are more likely to launch a new application than *Responsives*.

Joint analysis of application and user clusters

For each user cluster we determined the popularity of different application clusters (Table 2). In other words, we measure the association between users' overall phone revisitation patterns and the types of applications that they use. A Chi-Square test showed a significant association between User cluster and App cluster ($\chi^2(10) = 3480$, p < .01). While this statistical test checks for a two-way association (instead of one-way causality) between the two factors, we assume that although apps can be targeted at particular types of users and for use in particular contexts, users ultimately exercise control over which applications

App Clusters	User Cluster 1 U1	User Cluster 2 U2	User Cluster 3 U3
F1	768	1591	61
F2	4402	1686	100
M1	7346	9467	400
S1	635	1340	32
S2	2328	3305	390
H1	22747	15648	1313

Table 2. Popularity of application clusters for each user cluster

Regular expression	Example	Strategy	Instances	Popularity	Mean length
^[F]+$	FFFFFF	Initiate	3762	57%	2.27
^[F]+[B]+$	FFFBBB	Initiate-Revisit	1762	26%	4.17
^[F]+[B]+[F]+$	FFBBBFF	Initiate-Revisit-Conclude	598	9%	6.36
n/a	FFBFFB	Other	510	8%	8.50

Table 3. Summary of the 3 strategies we have identified for in-session backtracking. Mean length refers to the average number of application launches within a session, and is further elaborated in Table 4

they use. Therefore, we interpret our results in a one-way manner, hence revealing user preferences.

Analysis of in-session revisitation

Finally, our analysis looked at in-session revisitation patterns, as reflected in the text strings of F's and B's we constructed and analysed (Table 3). We calculated the revisitation ratio (ratio of B's to the size of the session string) as a function of the length of the string itself. The results indicated that as the size of the session increases then the likelihood of revisiting an application increases too, but eventually levels off at approximately 55% (n.b. revisitation can account for more than 50% of a session where individual apps are re-visited more than once).

Our results reveal a much higher level of switching back and forth between applications than previous studies. For example, Leiva *et al.* [17] reported that app switching was uncommon – switches to other apps interrupted 8% of app interactions.

Next, we manually analysed all session strings looking for patterns and similarities across multiple users. A qualitative analysis with two independent reviewers identified 3 distinct patterns which we then coded as regular expressions. The results suggest that 92% of all sessions fall into one of these 3 "Strategies" shown in Table 3, while only 8% could not be classified using our pattern-matching approach. Details about the mean length of different phases (in terms of number of application launches) are shown in Table 4.

Strategy	Length of Initiate	Length of Revisit	Length of Conclude	Mean length
Initiate	2.27	-	-	2.27
Initiate-Revisit	2.36	1.81	-	4.17
Initiate-Revisit-Conclude	2.50	2.53	1.33	6.36
Mean length	2.32	1.99	1.33	3.63

Table 4. For each phase of each strategy we show the mean length (number of application launches)

DISCUSSION

Micro-level analysis of smartphone revisitation patterns

Our analysis of in-session backtracking behaviour has enabled us to develop descriptive statistics of micro-level app usage behaviour on smartphones. Based on our results in Tables 3 and 4, we are able to define IRC (Initiate-Revisit-Conclude), a rudimentary model of usage. The model suggests that smartphone use – on a micro level – begins with an Initiation phase. This consists of launching

sequentially a number of distinct applications. This phase alone accounts for 57% of smartphone use (Table 3), and has a mean length of 2.27 apps (Table 4). Such an initiation strategy has been noted before [3], suggesting that certain apps are likely to encourage a user to engage with interaction with the smartphone.

However, we also find that there is a 35% chance that the Initiation phase will be followed by a Revisiting phase with mean length varying between 1.81 - 2.53 applications (Table 4). During this phase users will switch between applications they have launched moments ago since the session begun. This phase is indicative of users going back to the same applications possibly to check for new content, to confirm or verify some information, or for the sake of closure.

Finally, we observed that 9% of sessions contain a short (1.33 apps) Concluding phase. This consists of visiting a typically smaller number of new apps that were not launched during Initiation (and not visited during Revisitation). Our analysis suggests the existence of a user habit of checking one or two applications before turning off their phone. This mirrors the initiation habit reported previously [3], and is akin to a closing "ritual" for users. It is uncertain whether users choose to launch certain apps before switching off their smartphone, or whether launching certain apps leads them to switch off their smartphone.

Our rudimentary model can be used to make testable predictions. First, we expect that about 92% of all usage sessions can be described by this Initiate-Revisit-Conclude pattern. The remaining 8% may be attributed to users' inherent diversity. We have also identified certain interesting relationships between the different phases of our IRC model. The model predicts that in the presence of the Conclude phase we expect the Revisit phase to be longer than the Initiate phase (Table 4). However, in the absence of the Conclude phase then Revisit is actually shorter than Initiate. Thus, the model predicts that as Revisit grows in length it is increasingly likely that the user exhibits a Concluding phase before turning off the smartphone. Coupled with our result which poses an upper limit of approximately 55% on in-session backtracking – we can predict that as the length of Revisit approaches the length of Initiate, a Concluding phase is likely to take place and wrap-up the session.

Smartphone vs. web revisitation

Our analysis identified 6 clusters of applications, which we characterised as Fast, Medium, Slow, and Hybrid (Table 1). Surprisingly, our results bear interesting similarities and

differences with previous work on web revisitation patterns. Specifically, our analysis reveals the same four high level revisitation patterns as those reported in [1], which reported the habits of more than 600,000 users. We highlight these in Table 5, where we show for each type of revisitation pattern which applications we found and which websites were reported by [1].

Fast Revisitation on smartphones and the web
The websites with fast revisitations are fast monitoring and auto refresh pages, whose content changes at a fast rate. These bear great resemblance to social media and instant messaging applications found on smartphones in our study. We note that the Adar *et al.* study [1] does not mention social media or social networking sites, since it was conducted in 2008, and such websites have significantly grown in popularity in subsequent years.

Nevertheless, it is interesting to also note that Hub & Spoke websites have similarly fast revisitation patterns as social media and browsers on smartphones. Adar *et al.* reported that many of the pages in this fast cluster appeared to exhibit a hub-and-spoke revisitation pattern. i.e. a page containing many links to other pages, which the user is likely to visit briefly before returning to the original page to explore. An example would be a shopping site with a list of many items. Our analysis suggests that social media applications (and their respective instant messengers) and browser applications are being used in much the same way. Social media hubs also tend to "link" to external applications, taking the user to the browser, Youtube, music player and so on. Similarly, browser applications can also trigger other applications on smartphone.

Finally, it is important to highlight a peculiar difference between our findings and those in [1]. The fast revisitations in the Adar *et al.*'s web analysis have almost no long-term revisitation. On the other hand, our clusters of fast revisitation have non-trivial long term revisitation. We

attribute this difference to the fact that while there is a limited number of applications on one's phone, there is practically an unlimited number of websites that one can visit. Thus, while users tend to revisit certain websites frequently until they completely forget about them; on smartphones users tend to use certain applications frequently, then forget them, and then come back to them eventually. This can be attributed to the visual presence of applications in a phone's dashboard that act as a reminder and prominent anchor.

Medium Revisitation on smartphones and the web
Across the web and smartphones, we find that communication-related pages and applications tend to exhibit medium revisitation patterns. On the web, we find that webmail and forums are in this cluster, while in our data this refers to email and phone communication. Adar *et al.* have suggested that such medium revisitation patterns (1-4 hrs, or daily) are likely due to the timescale of human-to-human interaction using web mediated communication. In other words, this is driven by humans' tendency to communicate more on an hourly/daily basis rather than significantly faster or slower.

On smartphones we find a similar revisitation curve for phone and mobile-based email communication, suggesting that these communication technologies capture the same fundamental communication habits. We also note that communication with these technologies (asynchronous email and synchronous voice) differ from near-synchronous 'instant messaging' in our data, suggesting that these are appropriated differently by users, and serve a different purpose on smartphones.

An important difference we find between medium revisitation patterns on the web and on smartphones is attributed to home pages. Popular homepages appear in the medium revisitation cluster for web browsing, and they typically act as starting points for navigation in each browsing session. The equivalent in our dataset may be the UI or a popular application launcher, however we decided to filter these applications from our analysis due to their peculiar functionality. An informal assessment of their data has shown us that they would most likely fall in the fast clusters: the UI on smartphones acts both as a home page and a directory for finding applications, and therefore is more akin to a hub-and-spoke functionality.

Slow Revisitation on smartphones and the web
Once again, we find substantial similarities in slow revisitation patterns on the web and smartphones. In this cluster we find web pages with weekend activity and child-oriented content, which exhibit slow revisitation curves. These are conceptually very similar to the dating, health, fitness and game apps on smartphones which also exhibit slow revisitation. Thus, it can be argued that both categories reflect possibly individual or personal activities, which tend to follow a similarly slow periodicity on the web and smartphones. A further similarity we observe is that

Cluster Group	Centroid Curve	Description	Corresponding cluster group descriptions from Adar *et al.* [1]
Fast (F1,F2)		Instant Messaging, Browser, Social Media	Hub & Spoke, Shopping & Reference, Auto refresh, Fast monitoring, Pornography & Spam.
Medium (M1)		Email and Phone Communication	Popular homepages, Communication, .edu domain, browser homepages.
Slow (S1,S2)		Utilities, Multimedia, Health and Fitness, Games, Dating, Phone Settings	Entry pages, Weekend activity, Search engines used for Revisitation, Child-oriented content, Software updates
Hybrid (H1)		Documents, Notes, Video, Satnav	Popular but infrequently used, Entertainment & Hobbies, Combined Fast & Slow.

Table 5. For different revisitation patterns we show which smartphone apps we report, and which websites are reported by Adar et al. [1].

software update websites have a slow revisitation curve, as does the settings application and utilities on smartphones. These involve tweaking the configuration of one's system, and we find that both exhibit slow periodicity in both studies.

Do smartphones induce usage habits?

To investigate whether smartphones are a "special" technology that induces usage habits, we review the variety of evidence that our own study and previous studies offer. While we cannot explicitly test this assertion, we are able to interpret a variety of evidence across different studies. Our analysis has looked at temporal patterns of application usage at the macro level (revisitation patterns) and at the micro level (backtracking).

In our analysis we set out to investigate whether smartphones induce temporal usage habits that we do not see in other technologies. However, our analysis has shown that many of the macro-level behaviours we observed in our study bear close resemblance to web browsing habits. Hence our study provides some evidence to suggest that temporal patterns of usage behaviour on the web and on smartphones are driven by the nature of the service and information, and less so by the technology. The temporal behaviours that we have identified as common across the web and smartphones can be summarised as follows: (i) Users make quick revisits to applications and websites that contain fast-changing content or hub-and-spoke functionality. This is effectively the "checking" habit and "micro-usage" that has been reported on smartphones [13,25]. (ii) Websites and applications that facilitate asynchronous non-verbal or synchronous verbal human-to-human communication following medium revisitation. This is possibly driven by humans' tendency to communicate on an hourly/daily basis rather than significantly faster or slower [1]. (iii) Websites and applications that relate to personal activities follow a slow revisitation pattern. (iv) Websites and applications that involve technical tweaking and reconfiguration follow a slow revisitation pattern.

The first similarity we noted is the "checking" habit that leads to individuals checking and rechecking their phone. We argue that this habit is not unique to smartphones, as indeed users exhibit this habit on the web, and for similar purposes. Therefore, this habit is likely to be driven by the type of service or information that the use gets – irrespective of technology. In fact, the extent of similarities between user habits on the web and smartphones is such that we were unable to identify a temporal behaviour that is unique to smartphones. Most habits we identified on smartphones can be attributed to the service and not the technology: smartphones appear to be a conduit. Hence, we argue that, when viewed through the lens of revisitation analysis, smartphones do not induce usage habits at a macro-level.

However, we have found very interesting backtracking patterns which suggest that smartphones do induce usage

habits at the micro level. Surprisingly, our very simplified model of Initiate-Revisit-Conclude accounts for 92% of all usage we observed. Such a model has not been reported in web browsing studies. It is likely that the behaviour predicted by our model is induced by a limitation of smartphone technology: we have a relatively small number of applications available on our smartphone at any given time. Unlike web browsing where an almost unlimited number of pages is available for visit, on smartphones we have a few applications to choose from – and we tend to use and re-use them within individual sessions.

User characteristics or preferences?

To maintain our assessment that smartphones do not induce macro-level habits, we need to explain the substantial differences we observed between the different user clusters. Our analysis identified 3 clusters of users, each with different habits of application usage. Analysis of the adjusted residuals for Table 2 shows some interesting differences: Checkers (U1) use a disproportionately large amount of F1 and F2 (fast apps), and use disproportionately small amount of other apps. Waiters (U2) use disproportionately fewer F1 and F2, and relatively more slow, medium and hybrid apps. Responsives (U3) seem to use a mixture of slow and fast apps, as opposed to using mainly hybrid apps.

We contend that these are not habits induced by smartphones, but are more likely to reflect user preferences. Our justification for this argument is that we note that every user cluster had available applications from every application cluster. Therefore, each user cluster had the *potential* to use each application cluster, but did not choose to do so. Given our one-way assumption regarding causality between users and apps, we can infer that the skewed results in Table 2 are due to user preferences and not applications' availability. In other words, user preferences are driving the installation of applications on smartphones, and subsequently the nature of these applications is giving rise to the revisitation patterns we have observed in our users.

The mobile nature of smartphones, the fact that we carry them constantly [9], and their always-on connectivity all point to a very intuitive assertion in our discipline: people are induced to check their phones all the time. Yet our analysis has revealed both medium and slow revisitation patterns on smartphones that contradict this intuition. In addition, we show that smartphone revisitation patterns are very similar to desktop-based web browsing, which is a technology with rather different affordances. Hence, our findings offer evidence of certain innate habits that we exhibit across many technologies we use.

There is a way to bridge our discipline's intuitive understanding about smartphones with the empirical findings we have presented. If we consider smartphones not as hardware devices, but as sets of functionalities and services (i.e. apps), then we can make this reconciliation.

We suggest shifting our assertion from "checking our smartphones because they are mobile and they are available" to "checking our smartphones because of their apps." The mobility that our smartphones offer apparently makes no difference to our revisitation patterns. Whilst today we may be spending more time interacting with our smartphones than with our desktop computers a decade ago, in fact we are doing more of the same: we still exhibit fast, medium, and slow revisitation patterns. We still need to check for fast-changing content often, we still communicate with others at a medium pace, and we still engage in personal activities (and system tweaking) on a longer periodicity.

This brings us to the theoretical gap we pointed out in our motivation for this work, which in many ways remains elusive for our discipline. What are the fundamental forces that determine our use of mobile technology? Theoretical models have attempted to offer some explanation into how humans skew their behaviour when allocating time to applications, channels, searches, or web pages. A key purpose of technology use is to keeping informed, enhance social interaction, as well as simple entertainment [22,30], and people use technology that they believe will offer the gratifications they seek [28]. Therefore, as a ubiquitously available solution, mobile phones naturally compete against other information sources for users' attention. A useful analogy may be the use of the remote control, which has facilitated consumers to be more selective with TV channels so as to gain more gratification [28]. Do mobile phones make us more selective in our information needs? Our results suggest that different user clusters are more attracted to different types of services. However, additional investigation is necessary to address this question.

Further theoretical analysis of humans' skewed distribution of time allocated to activities has also been associated with people's ability to perceive their own past rate of activity [34] which subsequently leads them to accelerate or decelerate their rate of activity [4]. It has been hypothesized that this behaviour may be a fundamental evolutionary mechanism because it can synchronize populations of interacting species [5], stabilize them [23], and diversify gene pools [6]. In this sense, temporal skewing and backtracking may be far from confined to smartphones; rather it could be an instance of a broader human mechanism in how we allocate time to activities and how we regulate the rate of our activities.

LIMITATIONS
We note that our network monitoring application is likely to have attracted a specific type of user, which places some limitations on the generalizability of our findings. Due to our use of the Google Play App store to collect data we were unable to collect detailed information about the users included in our sample (e.g. gender, age, etc.).

We also note that although our study follows a "typical" [25] method to examine user habits, using quantitative data

to identify frequent behaviours (e.g. as in [27] [35]), other work has noted a distinction between frequent and habitual behaviour which suggests that the latter is a subset of the former. Oulasvirta *et al.* [25] suggests that habitual behaviours are extremely rapidly executed, whilst non-habitual behaviours are slower due to decision-making, and that habitual behaviours are also consistently associated with a particular triggering context. In our analysis we only consider temporal intervals between app visits, and the act of unlocking the phone as triggers for habitual behaviour. Therefore our analysis is unable to isolate habits that may be the result of other contextual cues (e.g. location), which may be unique to smartphone technology.

Although our revisitation analysis reveals repetitive behaviour, which is important for identifying smartphone habits, our analysis does not address differences in content contained within apps or websites. While we indicate that desktop browsing and smartphone app use exhibit similar characteristics, further research is required to understand whether there are content differences (e.g. more personal content) that distinguish smartphone and desktop use.

A further limitation of the analysis presented within this paper relates to the major advancements in desktop-based web browsing technology and web content that have taken place since the study reported in [1] was conducted. The significant growth in social media, and the proliferation of features such as tabbed web browsing, for example, make it possible that the revisitation patterns associated with desktop-based web browsing may look quite different today.

CONCLUSION
Motivated by prior work on analysis of revisitation patterns on the web, our paper has presented an analysis of the diverse ways that people revisit smartphone applications. Our revisitation analysis highlights two important findings. On a micro-level, we propose a simplified model of backtracking which accounts for 92% of usage on smartphones. On a macro-level, we find that smartphone revisitation bears remarkable resemblance to web browsing on desktops. This indicates that much of our habitual use of smartphones is not driven by the technology's characteristics, but rather by the characteristics of the services and information needs we have. Additionally, we call for researchers studying smartphone use to consider addressing the more fundamental driving forces that shape our use of smartphones, and indeed of technology in general. As a first step in this direction, in this paper we propose using revisitation analysis as a methodology for studying technology use.

ACKNOWLEDGMENTS

REFERENCES

1. Eytan Adar, Jaime Teevan, and Susan T. Dumais. 2008. "Large scale analysis of web revisitation patterns." In Proc. SIGCHI conference on Human Factors in Computing Systems, pp. 1197-1206.

2. Genevieve Bell and Paul Dourish. 2007. "Yesterday's tomorrows: notes on ubiquitous computing's dominant vision." Personal and Ubiquitous Computing 11, no. 2: 133-143.

3. 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 Proc. 13th international conference on Human computer interaction with mobile devices and services, pp. 47-56.

4. Dirk Brockmann, Lars Hufnagel, and Theo Geisel. 2006. "The scaling laws of human travel." Nature 439, no. 7075: 462-465.

5. James M. Bullock, Robert E. Kenward, and Rosie S. Hails. 2002. Dispersal ecology: 42nd symposium of the British ecological society. No. 42. Cambridge University Press.

6. Jean Clobert, Danchin, Etienne, André A. Dhondt, and James D. Nichols. 2001. Dispersal. Oxford: Oxford University Press.

7. Andy Cockburn and Bruce McKenzie. 2001. "What do Web users do? An empirical analysis of Web use." International Journal of human-computer studies 54, no. 6: 903-922.

8. Laura Dabbish, Gloria Mark, and Víctor M. González. 2011. "Why do i keep interrupting myself?: environment, habit and self-interruption." In Proc. SIGCHI Conference on Human Factors in Computing Systems, pp. 3127-3130.

9. Anind K. Dey, Katarzyna Wac, Denzil Ferreira, Kevin Tassini, Jin-Hyuk Hong, and Julian Ramos. 2011. "Getting closer: an empirical investigation of the proximity of user to their smart phones." In Proc. 13th international conference on Ubiquitous computing, pp. 163-172.

10. 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 Proc. 13th international conference on multimodal interfaces, pp. 353-360.

11. Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. "Diversity in smartphone usage." In Proc. 8th international conference on Mobile systems, applications, and services, pp. 179-194.

12. Denzil Ferreira, Eija Ferreira, Jorge Goncalves, Vassilis Kostakos, and Anind K. Dey. 2013. "Revisiting human-battery interaction with an interactive battery interface." In Proc. 2013 ACM international joint conference on Pervasive and ubiquitous computing, pp. 563-572.

13. Denzil Ferreira, Jorge Goncalves, Vassilis Kostakos, Louise Barkhuus, and Anind K. Dey. 2014. "Contextual experience sampling of mobile application micro-usage." In Proc. 16th international conference on Human-computer interaction with mobile devices & services, pp. 91-100.

14. Denzil Ferreira, Vassilis Kostakos, Alastair R. Beresford, Janne Lindqvist, and Anind K. Dey. 2015. "Securacy: An Empirical Investigation of Android Applications' Network Usage, Privacy and Security." In Proc. 8th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec).

15. Victor M. González and Gloria Mark. 2004. "Constant, constant, multi-tasking craziness: managing multiple working spheres." In Proc. SIGCHI conference on Human factors in computing systems, pp. 113-120.

16. Ke Huang, Chunhui Zhang, Xiaoxiao Ma, and Guanling Chen. 2012. "Predicting mobile application usage using contextual information." In Proc. 2012 ACM Conference on Ubiquitous Computing, pp. 1059-1065.

17. Luis Leiva, Matthias Böhmer, Sven Gehring, and Antonio Krüger. 2012. "Back to the app: the costs of mobile application interruptions." In Proc. 14th international conference on Human-computer interaction with mobile devices and services, pp. 291-294.

18. Philip Leroux, Klaas Roobroeck, Bart Dhoedt, Piet Demeester, and Filip De Turck. 2013. "Mobile application usage prediction through context-based learning." Journal of Ambient Intelligence and Smart Environments 5, no. 2: 213-235.

19. Yong Liu, Jorge Goncalves, Denzil Ferreira, Simo Hosio, and Vassilis Kostakos. 2014. "Identity crisis of ubicomp?: mapping 15 years of the field's development and paradigm change." In Proc. 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, pp. 75-86.

20. Yong Liu, Jorge Goncalves, Denzil Ferreira, Bei Xiao, Simo Hosio, and Vassilis Kostakos. 2014. "CHI 1994-2013: Mapping two decades of intellectual progress through co-word analysis." In Proc. 32nd annual ACM conference on Human factors in computing systems, pp. 3553-3562.

21. Gloria Mark, Daniela Gudith, and Ulrich Klocke. 2008. "The cost of interrupted work: more speed and stress." In Proc. SIGCHI conference on Human Factors in Computing Systems, pp. 107-110.

22. Denis McQuail. 1987. "Mass Communication Theory: An Introduction, Beverly Hills." Calif: Sage.

23. J. D. Murray. 1993. Mathematical biology. New York.

24. Hartmut Obendorf, Harald Weinreich, Eelco Herder, and Matthias Mayer. 2007. "Web page revisitation revisited: implications of a long-term click-stream study of browser usage." In Proc. SIGCHI conference on Human factors in computing systems, pp. 597-606.

25. Antti Oulasvirta, Tye Rattenbury, Lingyi Ma, and Eeva Raita. 2012. "Habits make smartphone use more pervasive." Personal and Ubiquitous Computing 16, no. 1: 105-114.

26. Antti Oulasvirta, Sakari Tamminen, Virpi Roto, and Jaana Kuorelahti. 2005. "Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile HCI." In Proc. SIGCHI conference on Human factors in computing systems, pp. 919-928.

27. Kurt Partridge and Philippe Golle. 2008. "On using existing time-use study data for ubiquitous computing applications." In Proc. 10th international conference on Ubiquitous computing, pp. 144-153.

28. Elizabeth M. Perse, 1998. "Implications of cognitive and affective involvement for channel changing." Journal of Communication 48, no. 3: 49-68.

29. Philipp Pushnyakov and Gleb Gusev. 2014. "User profiles based on revisitation times." In Proc. companion publication of the 23rd international conference on World wide web companion, pp. 359-360. International World Wide Web Conferences Steering Committee.

30. Quint Randle. 2003. "Gratification Niches of Monthly Print Magazines and the World Wide Web Among a Group of Special Interest Magazine Subscribers." Journal of Computer Mediated Communication 8, no. 4.

31. Choonsung Shin, Jin-Hyuk Hong, and Anind K. Dey. 2012. "Understanding and prediction of mobile application usage for smart phones." In Proc. 2012 ACM Conference on Ubiquitous Computing, pp. 173-182.

32. Linda Tauscher and Saul Greenberg. 1997. "How people revisit web pages: Empirical findings and implications for the design of history systems." International Journal of Human-Computer Studies 47, no. 1: 97-137.

33. Chad Tossell, Philip Kortum, Ahmad Rahmati, Clayton Shepard, and Lin Zhong. 2012. "Characterizing web use on smartphones." In Proc. SIGCHI Conference on Human Factors in Computing Systems, pp. 2769-2778.

34. Alexei Vazquez. 2007. "Impact of memory on human dynamics." Physica A: Statistical Mechanics and its Applications 373: 747-752.

35. Hannu Verkasalo. 2009. "Contextual patterns in mobile service usage." Personal and Ubiquitous Computing 13, no. 5: 331-342.

36. Tingxin Yan, David Chu, Deepak Ganesan, Aman Kansal, and Jie Liu. 2012. "Fast app launching for mobile devices using predictive user context." In Proc. 10th international conference on Mobile systems, applications, and services, pp. 113-126.

37. Chunhui Zhang, Xiang Ding, Guanling Chen, Ke Huang, Xiaoxiao Ma, and Bo Yan. 2013. "Nihao: A predictive smartphone application launcher." In Mobile computing, applications, and services, pp. 294-313. Springer Berlin Heidelberg.

38. Zhang, Haimo, and Shengdong Zhao. 2011. "Measuring web page revisitation in tabbed browsing." In Proc. SIGCHI Conference on Human Factors in Computing Systems, pp. 1831-1834.