

Understanding and Streamlining App Switching Experiences in Mobile Interaction

Alberto Monge Roffarello^{a,*}, Luigi De Russis^a

^a*Department of Control and Computer Engineering, Politecnico di Torino, 10129 Torino, Italy*

Abstract

Despite a large body of literature analyzing mobile device usage, app switching is still an overlooked interaction. To better understand and streamline the app switching experience in modern smartphones, we first explore how to automatically extract and characterize habitual app switching behaviors from smartphone usage data. By applying a data analytic methodology based on association rules to a large dataset of smartphone usage, in particular, we demonstrate that users repeatedly switch between the same applications under different contexts (e.g., location and time). We then implemented the methodology in RecApps, an interactive floating widget that proactively suggests the next apps to be used while the user is interacting with their smartphone. We evaluate RecApps through an in-the-wild study with 18 participants. Findings show that RecApps simplifies and supports the transitions between the users' favorite apps, while highlighting the need for novel interactions supporting app switching behavior. We use such results to explore trade-offs in the design space for proactively supporting app switching behavior in mobile interaction.

Keywords: Interaction types, Smartphones, App switching, App recommendations

*Corresponding author

Email addresses: alberto.monge@polito.it (Alberto Monge Roffarello),
luigi.derussis@polito.it (Luigi De Russis)

1. Introduction

In the last few years, smartphones have become ubiquitous devices [1], and their usage has increased dramatically [2]. One of the main reason for their success is their transformation from simple communication tools to computing devices offering a wide range of functionality accessible anywhere and every-
5 where [3]: through smartphones, users can nowadays read The New York Times while chatting with a friend on WhatsApp, they can watch films, check their emails, and browse social networks to pass the time.

Given the popularity of such devices, a large body of literature is dedicated
10 to study smartphone usage under different aspects. Previous works (e.g., [4, 5, 6]) analyzed smartphone usage datasets to characterize mobile sessions in terms of context, duration, frequency, and used apps. Other studies [2, 7, 8] related smartphones to habitual behaviors, by demonstrating that some types of smartphone usage, e.g., passively browsing social networks [8] or checking
15 emails [2], can become unconscious habits in the long term. Some previous works [6, 9] also analyzed the differences in the interaction with locked and unlocked smartphones, respectively.

In this vast literature, a particular kind of smartphone usage that is starting to emerge is the so-called *app switching behavior* [10], i.e., transitioning from one
20 app to another in the same usage session to consume content [11, 12]. Previous works already demonstrated the prevalence of this behavior: some researchers have found that the last used app is a strong predictor of the next app users are going to open [13, 14, 15], while others discovered that users transition from one app to another following established patterns, e.g., by typically switching from
25 mobile search engines to mobile apps [11] or by starting a usage session with a communication app [4]. However, such works only highlight high-level trends characterizing app switching behaviors, with data of different users that are mixed together to extract and analyze (typically off-line) global usage patterns. As a consequence, today’s smartphones offer limited support to transition from
30 one app to another. To close these gaps, this paper investigates the following

research questions:

RQ1. How can we automatically extract and characterize app switching behaviors, e.g., to understand which specific apps and contexts are involved?

RQ2. How can we design novel interactions to support these app switching behaviors in modern smartphones?

35

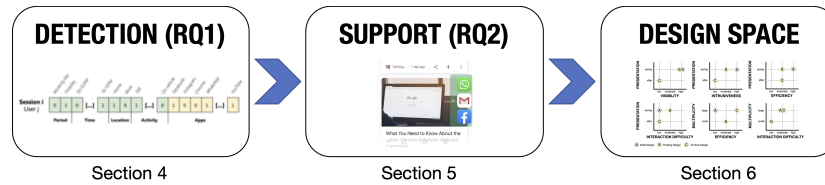


Figure 1: An overview of the different phases characterizing our work. We explored how to automatically *detect* app switching behaviors from smartphone usage data, we investigated how to *support* these behaviors in modern smartphones, and we derived a *design* space to proactively support app switching behaviors.

Figure 1 summarizes the different phases that characterized our work. To answer the first research question (*Detection* phase, **RQ1**), we explored a novel approach to automatically detect *habitual* app switching behaviors, i.e., transitions between apps that are repeated over time in different contexts, such as time of the day and location, from smartphone usage data. To this end, we adapted a recent data analytic methodology [16] based on association rules mining, and we demonstrated its applicability by exploiting a smartphone usage dataset collected in-the-wild from 46 users. By analyzing the extracted behaviors, we found that users often use two or more mobile apps in the same smartphone usage session by habitually switching between the same applications under stable contexts: habitual app switching, for example, is more common in the morning, when people wake up. Furthermore, app switching behaviors link different types of applications, and it is very common between web browsers, social networks, and messaging apps.

We used the outcomes of our first study to motivate and inform the exploration of our second research questions (*Support* phase, **RQ2**). The variety of

app switching behaviors that we extracted and analyzed, indeed, confirms that a promising (but underexplored) direction to improve the overall interaction with mobile devices would be to better support these behaviors on modern smart-
 55 phones [10]. While notifications and wallpaper widgets have expanded the ways people can interact with their smartphones, e.g., by enabling users to quickly check for new information of interest, today’s mobile user interfaces are only starting to explore novel interactions to streamline diverse device uses [6], and contemporary smartphones offer limited support to transition from one app to
 60 another. Despite the majority of today’s home-screen menu offers some type of app recommendations to find the next app to use, they often adopt very simple criteria, e.g., by showing the most recently used apps, only [17]. Moreover, users must explicitly stop to use the current app to see recommendations. Similarly, they had to perform a double tap on the home button (iOS) or a click on the
 65 app-list button (Android) to switch between already opened apps. To streamline app switching behavior and moving away from the “traditional” interactions for app switching, we thus present RecApps (*Recommending Apps*), an interactive floating widget designed to proactively support the transition from one app to another.

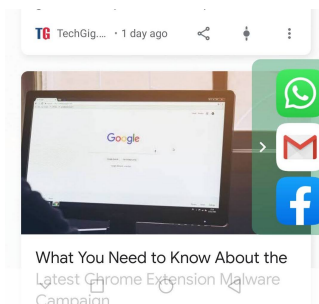


Figure 2: RecApps is an interactive floating widget for Android that proactively supports the transition from one app to another. When the user opens a mobile app, RecApps suggests which apps the user will probably use next through a floating widget. Under the hood, the app collects smartphone usage data and context information in background, and it extracts association rules that model habitual app switching behaviors.

70 RecApps dynamically implements our approach to automatically detect the

users' habitual app switching behaviors. It collects smartphone usage data and contextual information in background, and it periodically extracts association rules representing habitual app switching behaviors. The extracted rules are then used as suggested links: when the user opens an app that is involved in a switching behavior, a floating widget with a link to the apps that are typically related to the app in use appears on the screen (Figure 2).

We evaluated RecApps in a three-week long in-the-wild study with 18 smartphone users. Findings show evidence that proactively suggesting the next apps to be opened is a promising approach to streamline app switching behaviors: participants of the study liked the floating widget since the visualized suggestions speeded up their transitions from one app to another, especially when using social networks and messaging apps. RecApps did not change the overall behaviors of the participants with their smartphone, e.g., the average duration of their usage sessions and the number of apps they used. However, thanks to the availability of proactive shortcuts, it encouraged participants to transition between their favorite apps significantly more often. Despite the positive aspects, having an always-on widget on the right side of the screen sometimes interfered with participants' current activities, e.g., writing a message, and caused some unintentional clicks. Stemming from our findings, we present two alternative RecApps designs, and we discuss the design space for proactively supporting app switching behaviors by referring to different design dimensions, ranging from visibility to cognitive and physical demand (*Design* phase). Our exploration results in design recommendations on how to improve app switching interactions with mobile devices.

2. Related Work

2.1. Characterizing Mobile Device Use

Given the popularity of smartphones, there has been increasing research interest in studying mobile device usage under different aspects. Smartphones,

indeed, are nowadays used to perform a variety of tasks [3], ranging from enter-
100 tainment to productivity, in different contexts [18].

Several studies, e.g., [19, 4, 20, 21, 6, 5], analyzed smartphone usage sessions, i.e., users' interactions that start when the screen is turned on and end when the screen is turned off, in terms of overall duration, frequency, context, and involved mobile apps. Bohmer et al. [4] reported on the results of a large-scale analysis
105 on more than 4,000 smartphone users. They found that, on average, users spend roughly one minute with an app at a time, but such a time strongly depends on the app category. They also found that short sessions with only one app are very common. When users use more than one application in a session, instead, the first app is typically a communication tool. Another large-scale analysis of
110 smartphone usage based on detailed logs from 29,279 mobile phones has been recently published by Hintze et al. [5]. The authors demonstrated the influence of context, e.g., locations, on the frequency of mobile device usage. Furthermore, they highlighted usage differences between smartphones and tablets, showing that smartphones are used almost thrice as often as tablets, with usage sessions
115 on tablets that are typically longer. Also other works highlighted the influence of users' locations on mobile device use [19], while others related differences in smartphone usage to social context [20]. Banovic et al. [6], instead, focused on the time duration of mobile device use, by classifying usage sessions into glance, review, and engage sessions. In their work, in particular, the authors
120 developed ProactiveTask, an interactive lock-screen tool providing short tasks to users to streamline review sessions, i.e., brief user interactions with one or more applications. As in the work of Banovic et al., we present a tool to streamline a particular type of user's interaction with mobile devices, i.e., app switching behaviors.

125 Besides the presented analysis, researchers have also demonstrated that smartphone usage is governed by different kinds of cognitive [22] and psychological factors [23]. In the last few year, smartphone usage is so increased [2] that it has even been related to overuse problems [24] and addictive behaviors [25]. Scientific evidence shows that there are substantial individual differences in how

130 users may engage with their personal device [1, 6, 12], e.g., given the variety
of available mobile apps. However, some previous works clearly demonstrate
that, from an abstract point of view, different usage patterns are shared across
people. Several studies, for instance, related smartphone usage to habitual be-
haviors [26, 8, 27], i.e., recurrent phone usage sessions that are repeated under
135 stable contextual cues [2]. Oulasvirta et al. [2], for instance, analyzed checking
habits, i.e., “brief, repetitive inspection of a dynamic content quickly accessi-
ble on the device.” What makes users continuously check their smartphones
until their needs are satisfied [28] is often a self-interruption to check online
contents [29], missed calls, or messages [23]. Shin et al. [30] found that users
140 with self-identified problems in smartphone usage are more susceptible to check-
ing habits, and they tend to use many different apps during the same phone
usage sessions. By analyzing 3 months of application launch logs through a
methodology extensively adopted in the context of web browsing, i.e., revisita-
tion analysis [31], Jones et al. [11] discovered that much of our habitual use of
145 smartphones is not driven by the technology’s characteristics, but rather by the
characteristics of the services we have. On smartphones, in particular, we have
a few installed applications to choose from, and we tend to use and re-use them
within individual sessions. Such revisitation patterns, also called app switching
behaviors [10], are still an overlooked aspect of mobile device usage, and they
150 are not proactively supported by modern smartphones, yet. To transition from
an app to another, indeed, users must explicitly stop the use of the current app,
e.g., by performing a double tap on the home button, and look for the new one to
be opened, e.g., by browsing the list of apps that are currently opened. Conse-
quently, smartphones are often unable to provide information at an appropriate
155 time and quickly enough for the interaction to be worthwhile [18].

Turner et al. [10] recently confirmed the need to support app switching
behaviors by showing that there is underlying similarity in such patterns across
users. In this paper, we further analyze the habitual nature of app switching
behaviors by applying a data analytic methodology based on association rules
160 mining over a large dataset of mobile device usage. Furthermore, we explore

how to support recurrent transitions between apps by presenting RecApps, an interactive floating widget that proactively suggests the next apps to be used.

2.2. Predicting the Next App

RecApps dynamically provides users with links to easily switch between mobile applications, with the aim of streamlining habitual app switching behaviors. Therefore, our work is also related to the body of research on app recommendations. The problem of predicting which app users are going to use has been already investigated by other researchers in the past. Being able to perform such a prediction might produce different advantages, e.g., pre-loading the right app to improve memory management and app execution [21, 13, 32], or to highlight desired apps in the home screen for quicker launches [13, 33]. Sun et al. [33], for instance, developed AppRush, an adaptive home screen widget presenting shortcuts towards mobile apps most likely to be launched at a certain time, based on the user’s app usage history. Instead of highlighting apps in home screens, in our work we dynamically visualize links to other mobile applications on top of the app that is currently used. Furthermore, we explicitly focus on *habitual* connections between different apps under multiple contexts and by continuously adapting the suggested shortcuts to the current usage session.

Regardless of the underlying goal, different algorithms for app prediction have been proposed in the last years. Yan et al. [21] presented FALCON, a system that exploits contexts such as user location and time of the day to predict app launches before they occur. The goal is to provide support for pre-launching applications, thus reducing the perceived delay in opening them. With a similar goal, Parate et al. [32] designed an app prediction algorithm that can work without prior training. The algorithm is able to predict which app will be used and when, without requiring additional sensor context. Zou et al. [34] explored light-weighted Bayesian methods to predict the next app based on the app usage history. Huang et al. [13] used the Nokia MDC dataset to test a linear and a Bayesian model for app prediction, by exploiting different contextual features such as time, location, and last used app. They found that

the last used app is a strong contribution to the prediction accuracy, with the linear model that turned to be the more effective in combining all the contextual information.

Other studies found a correlation between the last used app and the next app users are going to use [14, 15]. Leroux et al. [14] presented a framework for the prediction of a user’s future mobile application usage behavior. The framework has been developed to continuously monitor several contextual information on the smartphone, e.g., locations and used apps, with the aim of automatically deducing usage patterns. Baeza-Yates et al. [15] proposed a prediction mechanism to show users which app they are going to use in the near future. They modeled app prediction as a classification problem, and they exploited two kinds of features, i.e., basic features obtained from smartphones sensors and session features modeling sequential patterns of app usage. In their experiments, the authors demonstrated that session features are effective to boost the performances of their algorithm.

The fact that the last used app is a strong predictor of the next app suggests the need of further exploring app switching behaviors. The majority of the aforementioned works, however, often exploits *representational contexts* that are usually defined before the user interacts [35]. Besides taking into account past usage data and context, in our RecApps tool we also focus on the *current* phone usage session, with the aim of making personalized and dynamic recommendations to a user that is currently engaged with her smartphone. We were inspired by the *interactional context* concept defined by Natarajan et al. [35]. According to the authors, app recommendations must be made available dynamically as the user interacts with the system. When users interact with a smartphone, indeed, they often click on items that are of interest in the current context. Finally, differently from the recommendation algorithms presented here, we are less interested in predicting the single, next used app with the highest accuracy. Rather, our approach is more generic, and aims at extracting connections between *sets* of apps, e.g., to model a user that opens WhatsApp after having used both the Camera and the Contacts app. Our main goal, in particular, is

to explore whether the RecApps approach, e.g., the floating widget, is effective to streamline app switching interaction. For this reason, while previous work on app prediction leverages off-line evaluations, mainly, our work includes an in-the-wild user study.

3. Detecting App Switching Behavior

We explored how to automatically extract and characterize app switching behaviors from smartphone usage data (**RQ1**) by applying a data analytic methodology based on association rules mining to a mobile device use dataset collected in a real-world setting. On the one hand, our goal was to further dig into the anatomy of app switching behavior. On the other hand, we laid the foundation for detecting and supporting app switching in real-time (see Section 4).

3.1. Automatic Detection of Habitual App Switching Behaviors

The first step to support users in transitioning between their favorite apps is to detect, preferably in real-time, what are their most common app switching behaviors. To this end, we adapted a data analytic methodology based on association rules mining [16] to extract *habitual* app switching behavior. The methodology was designed to extract, directly on a smartphone and in real-time, the smartphone habits of a user. A habit is defined as an association rule between a set of contextual cues (the rule’s antecedent) and the usage of one or more mobile apps (the rule’s consequent). Figure 3 shows two examples (A and B) of smartphone habits that can be extracted through the methodology. Contextual cues may include information like time or current location, but also the usage of a set of apps inside the session. The habit (A) describes a user that at work, between 10 and 12 AM, habitually uses Facebook and Instagram. The habit (B), instead, is a classic example of an app switching behavior: when the user is at home, they typically transition from WhatsApp to Twitter. We exploited the ability of the methodology to treat apps both as antecedent and consequent of association rules to extract switching behaviors that are repeated over time, i.e., links between the same mobile apps, in different contexts.



Figure 3: Two examples of smartphone habits that can be extracted through the applied methodology. The habit (B) is a habitual app switching behavior between WhatsApp and Twitter that happens when the user is at home.

To initially demonstrate the applicability of the methodology to the detection of app switching behaviors, we exploited a mobile usage dataset collected in-the-wild from 46 users (36 male and 10 female). Users were mostly students (44), and they were on average 22.76 years old ($SD = 2.31$). Data was collected through Socialize [16], an Android mobile application with which users can monitor their smartphone usage. Socialize silently collects different users' information and make it available, in an anonymous form, on a Firebase [36] dataset. Table 1 describes the information available in the dataset we used to explore habitual app switching behavior: we exploited phone-related information, i.e., screen and app events, and contextual information as well, i.e., activity and location events, in addition to time.

Following the methodology, we preprocessed smartphone usage data to build usage sessions for each user. First, we isolated pairs of consecutive lock-unlock screen events to delineate the start and the end of each usage session, i.e., the session-window. Then, we used each session-window to extract *a*) the mobile apps used during the session and *b*) the set of contextual information characterizing the session. As contextual information, in particular, we considered the *time* (e.g., 10-12 AM) and the *period* (i.e., *working day* or *holiday/weekend*) of the session, the physical *activity* and geographical *location* characterizing the session, if available, and the mobile *apps* used in the entire session. After the preprocessing step, each smartphone usage session is transformed in a transactional data format (Figure 4).

Table 1: The information of the mobile usage dataset [16] we used to explore habitual app switching behavior.

Information	Description
Activity Events	Start/stop a given activity, i.e., <i>still</i> , <i>walking</i> , <i>running</i> , <i>cycling</i> , and <i>on vehicle</i> . Each activity event includes a timestamp, an activity, and the type of the event, i.e., start or stop.
Location Events	Enter/exit a given location area. The application forces users to define at least their <i>home</i> and <i>work</i> locations. Each location event includes a timestamp, a location, and the type of the event, i.e., enter or exit.
Screen Events	Lock/unlock the smartphone screen. Each screen event includes a timestamp and the type of the event, i.e., lock or unlock.
App Events	Open/close a given mobile app. Each app event includes a timestamp and the type of the event, i.e., open or close.

The resulting transactional dataset is finally used to mine association rules, i.e., smartphone habits, through the Apriori algorithm [37]. For each user, promising association rules can be filtered by using common metrics for association rules evaluation [38], e.g., support, confidence, and lift. We enforced support greater than 1% and confidence greater than 51% to prune uncorrelated combinations. Furthermore, we discarded association rules with lift less than or equal to 1 to prune negatively correlated combinations. Finally, we excluded all the rules that did not represent a habitual app switching behaviors, i.e., rules without any mobile apps in their antecedent, as in Figure 3 (A).

3.2. Results

3.2.1. Phone Usage Sessions Overview

By preprocessing the exploited dataset, we obtained a total of 133,988 smartphone usage sessions with a median duration across all sessions of 56.73 seconds. Note that we use the median since the duration of smartphone usage sessions is typically characterized by a long tail distribution [6]. Of these sessions, 30,524

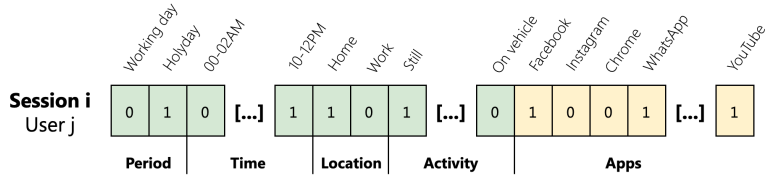


Figure 4: An example of a smartphone usage session represented in a transactional format. Besides mobile applications, the vector involves contextual information, e.g., activity and location. In each entry, in particular, the value 1 means the presence of the item in the corresponding transactional session, and the value 0 represents the absence of an item in that session.

(22.78%) were associated to a geographical location, and 132,489 (98.88%) to a performed physical activity. The time intervals during which users produced more sessions was between 6 and 8 PM (12.81%) and between 8 and 10 PM (12.80%). Not surprisingly, 97,481 sessions (72.75%) happened during business days, while the remaining 36,507 (27.25%) happened on a holiday/weekend.

Figure 5 characterizes the number of mobile apps included in the collected smartphone usage sessions. We found that 11,611 sessions (8.68%) did not include any mobile application. These sessions are what Banovic et al. [6] call “glance sessions”: brief interactions lasting some seconds to check information, e.g., the current time, on the lock screen. As in the work of Banovic et al., their median duration was roughly 13 seconds (13.92).

Other 67,314 sessions (50.24%) included the usage of a single mobile app, only, with a median duration of 36.83 seconds. In the categorization proposed by Banovic et al. [6], these sessions can be classified as “review sessions,” i.e., interactions lasting less than 60 seconds to consume content and provide quick input to an application, e.g., opening Gmail and briefly scrolling through existing emails.

The number of mobile apps used in the remaining 55,063 smartphone usage sessions was on average 2.89 ($SD = 1.39$), and ranged from 2 (30,226 sessions) to 20 (1 session). This means that a considerable amount of mobile device use (41.09%) included at least a transition between two applications. These

“engage sessions [6]” typically last more than 60 seconds. In our case, the
310 median duration was 149.63 seconds.

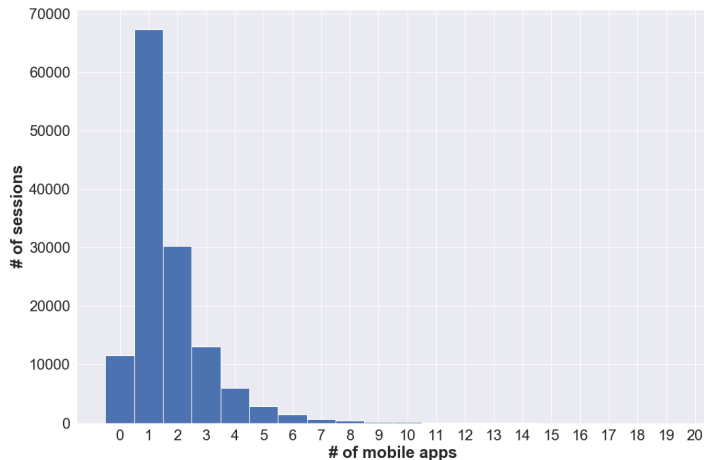


Figure 5: Histogram of the number of mobile apps included in the collected smartphone usage sessions.

3.2.2. Habitual App Switching Behaviors

From the 55,063 smartphone usage sessions involving more than one app, we extracted a total of 16,227 habitual app switching behaviors. This suggests that such habits characterize a small, specific portion of the overall usage of
315 smartphones. On average, each user demonstrated 352.76 switching habits, with a standard deviation of 740.40 habits. These relatively high numbers are due to the fact that the exploited methodology treats each calculated association rule as a separate habit. It is therefore possible to have very similar rules that differ by a single contextual cue, with the same app switching behavior (e.g.,
320 WhatsApp to Facebook) that can be specified in different “granularities” and for different contexts. The same behavior, for example, may be particularly common in the morning, but also in general during all day. What is clear from this analysis is that the number of such behaviors differed significantly across users, and confirms that there are non-neglectable differences in how people

325 interact with their mobile devices [39].

Figure 6 shows the distribution of the number of apps included in the extracted habitual app switching behaviors. The figure highlights different kinds of app switching behaviors. While association rules modeling a direct transition between two apps were very common (see the first two bars of Figure 6), we also found several association rules with *multiple apps* in their antecedents. This means that some combinations of mobile apps can spur users to open other applications. An example we found during our analysis was a user that habitually used WhatsApp when using both the Camera and the Contacts apps, probably to share her photos with her friends.

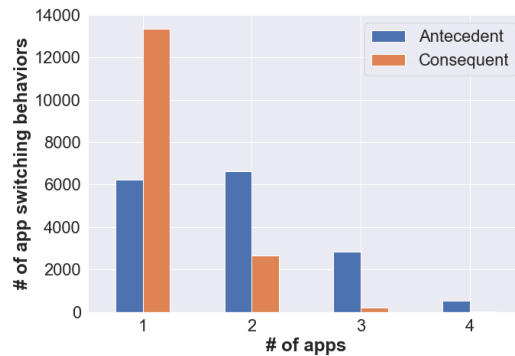


Figure 6: Distribution of the number of apps included in the antecedent and consequent of the extracted habitual app switching behaviors.

335 Besides the different number of involved apps, the extracted habitual app switching behaviors were characterized by different contextual cues. Users, for instance, habitually transitioned between different apps at home (4,334 rules), during business days (5,695 rules), and when they were sitting in a place (7,003 rules). Interestingly, while the majority of the sessions in the exploited dataset happened in the evening, we found that the time interval during which habitual app switching behaviors were more common was between 6 and 8 AM (628 rules). This suggests that users are likely to routinely transition between the same mobile applications when they wake up in the morning.

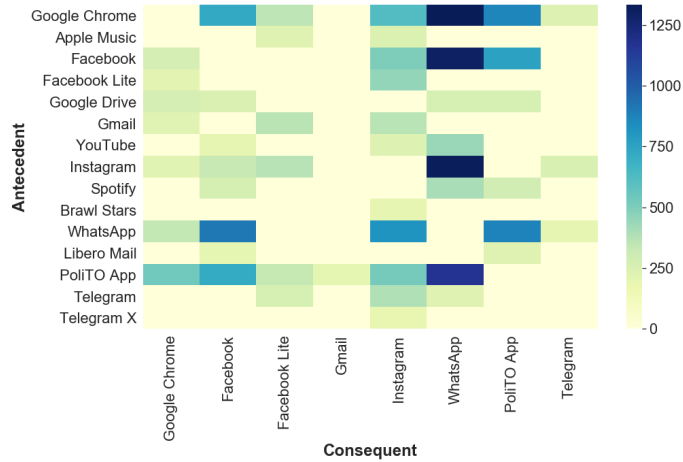


Figure 7: An heatmap representing the extracted habitual app switching behaviors on the basis of the involved mobile apps.

Finally, the heat map in Figure 7 describes the extracted association rules on the basis of the mobile apps involved in their antecedents and consequents, respectively. While the set of used apps varied significantly across participants, we were able to find common mobile applications that were linked in several app switching behaviors. The most common links, in particular, involved a messaging app (e.g., WhatsApp), a web browser (e.g., Chrome), social networks (e.g., Facebook and Instagram), and an educational app (e.g., PoliTO App). We found, for instance, 1,334 rules containing Chrome in the antecedent and WhatsApp in the consequent, and 914 rule containing WhatsApp in the antecedent and Facebook in the consequent. Although the analyzed data do not allow to understand the underlying usage motivations, such common app switching behaviors suggest that there could be several reasons for which users habitually switched between these apps, from completing a task that involves multiple apps to satisfy an established usage routine. The usage of a browser like Chrome before a messaging app like WhatsApp, for example, may suggest that the user has shared some content found online on one of her chats. Using Facebook after WhatsApp, instead, is probably related to a pure usage habit,

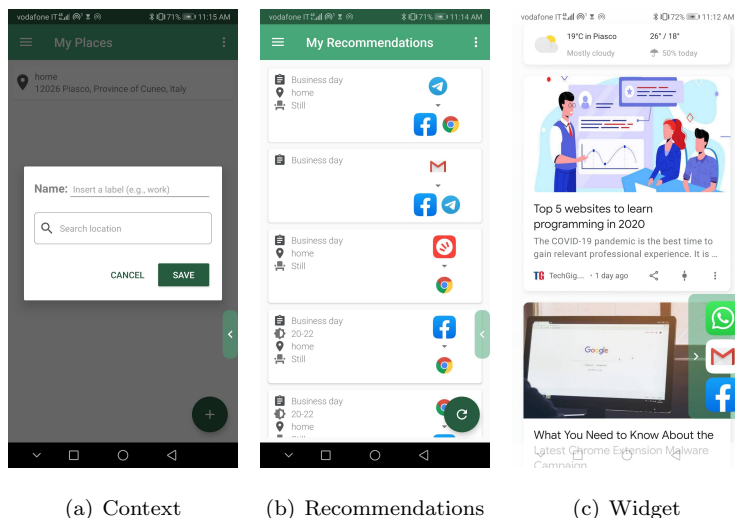
e.g., checking social networks after having received (and perhaps answered to) a message on WhatsApp.

4. Supporting App Switching Through RecApps

Our first study shows evidence that users habitually switch between different mobile apps in different contexts, and that the applied data analytic methodology is able to automatically detect these behaviors from smartphone usage data. As contemporary mobile devices offer limited support to transition from one app to another, the second part of our research focused on how to streamline such an interaction (**RQ2**). To this end, we implemented the data analytic methodology in RecApps, an interactive floating widget whose goal is to allow users to switch to the next app without the need of closing the app that is currently used and browsing the home-screen menu or a list of recently used apps.

4.1. RecApps

We developed a first version of RecApps as an Android application implementing the same data analytic methodology described in Section 3. The app silently collects smartphone usage data in background, and, when available, it associates contextual information, i.e., time, period, physical activity, and current location, to each smartphone usage session. Activities, i.e., *still*, *walking*, *running*, *cycling*, and *on vehicle*, are detected through the Google Activity Recognition APIs [40]. Locations of interest, e.g., *home* or *workplace*, can be instead defined at the startup of the application (Figure 8(a)). Twice a day, RecApps uses the collected information to recalculate, directly on the smartphone, association rules modeling habitual app switching behaviors of the user. Similar rules, e.g., those that include the same contextual information, are then merged together to generate a set of “recommendations” that can be viewed at any time in a dedicated screen of the app (Figure 8(b)). In our work, we define recommendations as shortcuts that allow users to repeat their habitual



(a) Context (b) Recommendations (c) Widget

Figure 8: The RecApps application in a nutshell. Users can define one or more preferred locations to be monitored (a). RecApps collects usage and contextual information to extract association rules and generate recommendations (b). Recommendations are dynamically proposed to the user through a vertically-floating widget (c).

switching behaviors more quickly. The first recommendation of Figure 8(b), for instance, suggests that during a business day, when the user is sitting at home
 390 using Telegram, she typically switches to Facebook and Chrome.

Recommendations are proactively proposed to the user through a vertically-floating widget on the right side of the screen (Figure 8(c)). When the antecedent of a calculated association rule is verified, e.g., the user opens a given
 395 app in a specific contextual situation, a link to the apps that are typically related to the app currently in use appears on the screen. Such links are continuously adapted to the apps that the user is currently using in the session. When the user opens 2 different apps in the same session, for instance, RecApps updates the widget content by checking if there are association rules involving those 2
 400 apps as the rule’s antecedent.

When visualizing the widget, the user can simply click on a displayed icon to transition towards the related app. The number of apps displayed in the widget is not fixed and it is automatically computed by RecApps according to

the association rules that are currently active, i.e., those that reflect the current
405 contextual situation. The widget can be reduced at any time by the user (as in
Figure 8(a) and Figure 8(b)), and it can be vertically moved on the right area
of the screen.

4.2. User Study

We conducted a user study to evaluate RecApps in a real-world setting. We
410 investigated, in particular, how users would interact with RecApps, and whether
the proactive widgets influenced participants' behavior with smartphones. Our
aim was to define a design space for proactively supporting app switching be-
haviors.

4.2.1. Participants

We recruited a total of 18 participants (11 female and 7 male) by exploiting
415 mailing lists of different university courses and by sending private messages to
our social circles. On average, participants were 23.05 years old ($SD = 3.54$).
Except for 3 employees, they were mostly university (13) or high-school (2)
students.

4.2.2. Method

The study was conducted between July and August 2020, and lasted three
weeks. We uploaded RecApps on the Google Play Store, and we asked partici-
pants to install it. In the first week, RecApps ran in the background by silently
logging usage data, without displaying any widget (*baseline* phase). This en-
425 abled RecApps to get enough data to calculate the first recommendations. After
seven days, a notification alerted the participants that recommendations were
active. From that moment, RecApps started to periodically recalculate associa-
tion rules modeling habitual app switching behaviors, and participants could use
the displayed widgets to switch between their favorite apps (*RecApps* phase).
430 After three weeks, we conducted a brief semi-structured interview with partici-
pants to get their qualitative feedback. Due to the Covid-19 pandemic [41],
interviews were conducted remotely via the Zoom video conferencing tool [42].

In this final part of the study, we explored participants' experience using RecApps by taking advantage of the following open-ended questions:

- 435 • Overall, what has been your experience with RecApps? Did you like it? What were its advantages and disadvantages?
- Did you use the visualized recommendations to switch between apps? When did you find these suggestions more useful?
- Has RecApps changed the way you use the phone?
- 440 • Do you have any suggestions to improve RecApps?

4.2.3. Data Collection and Metrics

During the study, we collected anonymous data by using Cloud Firestore [43], a Google-powered NoSQL database for storing, syncing, and querying data of mobile and web apps. We collected usage data related to participants' smart-
445 phone usage sessions, including the included apps and the associated contextual information, e.g., current location. We averaged these measures for each participant, by computing the following metrics:

Duration. The average duration of the smartphone usage sessions performed by a user.

450 **Unique Apps.** The average number of unique apps that a user uses in her smartphone usage sessions.

App Openings. The average number of distinct app openings, including multiple openings of the same app, that a user performs in her smartphone usage sessions.

455 Furthermore, we kept track of the calculated recommendations. We measured, in particular, how many times participants saw a widget with available recommendations, and how many times they clicked on a recommended item to perform an app switching behavior. To measure the usefulness of RecApps, in particular, we computed the following metric, averaged for all the participants:

460 **Used Widgets (%)**. We considered a widget as “used” when participants
clicked on one of its suggested apps. The percentage of used widgets is
therefore computed as the number of used widgets over the total number
of time participants saw a widget.

4.3. Results

465 Results are organized as follow. First, we explore the frequency of use of
RecApps during the study. Then, we move to investigate the content of the
displayed widgets and its influence on the probability of a widget to be used.
Finally, we analyze the influence of RecApps on the participants’ behavior with
smartphones, we compare the adopted recommendation algorithm with basic
470 approaches suggesting the most used apps, only, and we report on the perceived
advantages and disadvantages of RecApps collected during the final interview.

4.3.1. Frequency of Use

In the 14 days during which recommendations were active, each participant
clicked on a suggested app to directly switch from the current app to another
475 71.61 times on average ($SD = 124.90$, $median = 30.50$). More specifically,
each day RecApps helped participants to simplify 7.58 habitual app switching
behaviors on average ($SD = 11.80$, $median = 3.00$). Overall, the visualized
widgets were actively used (**used widgets (%)** metric) in 5.80% of cases. In-
terestingly, as shown in Figure 9, participants clicked on suggested apps more
480 frequently as the day passed. Indeed, the percentage of used widgets increased
during the study and ranged from 4.64% of day 8 to 8.01% of day 21. A Pear-
son’s correlation test confirmed that there was a moderate positive correlation
between the day of the study and the percentage of used widgets ($r(19) = 0.608$,
 $p < .05$). Given that app switching behaviors are often executed out of a habit,
485 sometimes even unconsciously, this may suggest that it takes time for users to
abandon their existing method for switching between apps in favor of RecApps.
Another possible explanation, however, is that RecApps was able to continu-
ously improve the accuracy of the computed recommendations by collecting and

exploiting more data.

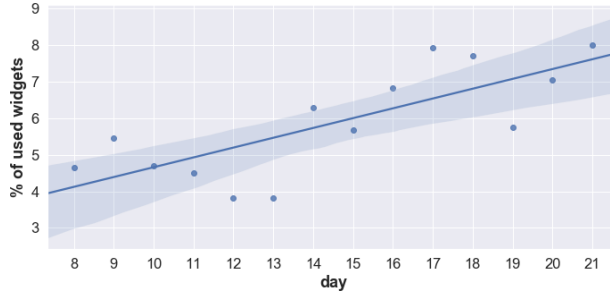


Figure 9: The percentage of used widgets, i.e., widgets for which participants clicked on a suggest app, during the study. As the days passed, participants clicked on a recommended app to switch to another app more frequently.

490 Such an ability of RecApps was confirmed by a participant in the final in-
terview:

*I installed RecApps when I was on vacation, so when I came home
recommendations did not reflect my habitual smartphone usage. How-
ever, I noticed that the app quickly adapted to my ‘return to everyday
495 life’.*” (P12)

While the reported numbers may appear low, they are perfectly in line with
the results described by the authors of ProactiveTasks [6], another attempt to
support proactive interactions on smartphones. As in ProactiveTasks, in par-
ticular, interaction with RecApps was not supposed to be participants’ primary
500 goal. Rather, RecApps provided an optional mechanism to simplify app switch-
ing behaviors that were habitually performed: as described in Section 3, such
habitual app switching behaviors characterize only a small portion of the overall
usage of smartphones.

4.3.2. Widget Content

505 In line with what we found in our formative study (Section 3.2.2), the most
common clicked recommendations referred to app switching behaviors involv-
ing social networks and messaging apps. Participants, for instance, clicked 337

times on the WhatsApp icon while using Instagram, and 224 times on the Instagram icon while using WhatsApp. The usefulness of RecApps to streamline such *social-to-chat* and *chat-to-social* switching behaviors was confirmed by participants in the final interview. P12, for instance, said “*I often used RecApps to switch from social networks to Telegram*”, while P2 said “*I really appreciated the visualized shortcuts while I was browsing social networks, I could move between Facebook and Instagram really fast.*” According to the participants, however, RecApps was also useful for other scenarios, e.g., opening the web browser when needed (“*it was useful to move from one social to another, but also to google something faster*”, P7) or connecting two interrelated apps (“*I liked the fact that I had a shortcut to the gallery when I was using the camera*”, P2).

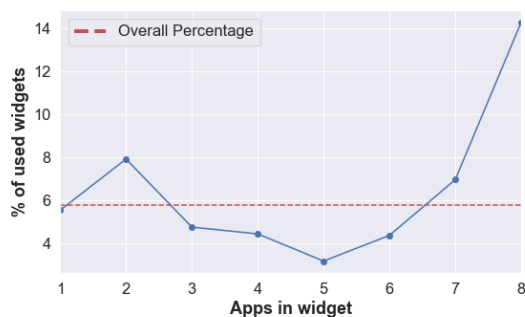


Figure 10: The number of apps displayed by the RecApps widgets was not fixed, but ranged from 1 to 8. The figure shows how the number of displayed icons influenced the probability of the widget to be used.

In RecApps, the number of apps displayed by the widgets was not fixed, but was automatically computed by RecApps on the basis of the currently active association rules. Such a number ranged from 1 (54.36% of all the widgets) to 8 (0.03% of all the widgets). As reported in Figure 10, we found that the number of icons displayed in a widget influenced the probability of the widget to be used, i.e., with a click on a displayed icon. The percentage of used widgets was similar or higher than the overall average (5.80%) when widgets were “small,” i.e., with 1 or 2 icons, respectively. Such a percentage decreased for bigger

widgets, i.e., those displaying 3 to 6 apps. By covering a relatively large part of the right-hand side of the smartphone screen, such widgets interfered with the user’s activities, and they were often closed without being inspected, as reported by P18: “sometimes the widget opened with so many icons...it was a bit invasive, I closed it right away.” Finally, very big widgets displaying 7 and 8 icons can be considered as outliers. Indeed, only 43 widgets with 7 icons appeared during the study across all users. This number was even smaller for widgets displaying 8 icons, that appeared 14 times, only. Furthermore, the high usage rate of such widgets can be associated to unintentional clicks. As some participants noted, recommendations were sometimes used with “unintentional touches” (P5), especially when there were several icons on the screen: “since widgets appear on top of the screen, sometimes you trigger a shortcut without really wanting it, especially when there are too many apps in the widget” (P11).

4.3.3. Influence on Smartphone Usage Patterns and Comparison with Baseline

To understand whether and how RecApps influenced participants’ behavior with smartphones, we compared some of the computed metrics before the widgets were available (*baseline* phase) and after the app started to visualize them (*RecApps* phase). We conducted a series of Wilcoxon Signed-Ranks tests on the **duration**, **unique apps**, and **app openings** metrics, respectively. The reason for choosing a non-parametric test that does not assume a normal data distribution was due to the high variability in the data under analysis. Indeed, participants used their smartphones in different ways. Two Shapiro–Wilk tests, in particular, revealed that usage sessions did not follow a normal distribution in terms of duration ($p < 0.05$) and number of included apps ($p < 0.05$).

Table 2 shows how the usage of RecApps influenced the average duration of the participants’ smartphone usage sessions. The median duration of a session in the *baseline* phase was 3.46 minutes (range = [1.23 -24.33]). The median duration was instead 2.72 minutes in the *RecApps* phase (range = [1.30 -399.36]). A Wilcoxon Signed-Ranks test did not reveal significant differences ($p > 0.05$).

Table 3 demonstrates that also the average number of unique apps in a

Table 2: How the usage of RecApps influenced the average **duration** of the participants’ smartphone usage sessions. A Wilcoxon Signed-Ranks test did not reveal significant differences.

	Baseline		RecApps		p
	Median	Range	Median	Range	
Duration [min]	3.46	[1.23 - 24.23]	2.72	[1.30 - 399.36]	0.263

Table 3: How the usage of RecApps influenced the **unique apps** and the **app openings** metrics, respectively. Two Wilcoxon Signed-Ranks tests did not reveal significant differences.

	Baseline		RecApps		p
	M	SD	M	SD	
Unique apps [#]	1.65	0.23	1.73	0.31	0.225
App openings [#]	7.26	3.26	14.13	29.58	0.329

session and the average number of app openings were not significantly influenced by the usage of RecApps. Results, however, highlight a possible trend towards sessions with more app openings: although not statistically significant ($p >$
560 0.05), in particular, the average number of unique apps in a session varied slightly from 1.65 (SD = 0.23) to 1.73 (SD = 0.31), and the number of app openings slightly increased from 7.26 (SD = 3.26) to 14.13 (SD = 29.58).

Besides investigating whether and how smartphone usage sessions changed after using RecApps, we also analyzed whether RecApps encouraged the repetition of the participants’ habitual app switching behaviors. To perform such
565 an analysis, we first extracted the most common app switching behaviors reproduced by the widgets displayed by RecApps:

- Instagram to WhatsApp and WhatsApp to Instagram were displayed in a total of 11,407 widgets;
- 570 • Facebook to Instagram and Instagram to Facebook were displayed in a total of 4,451 widgets;
- Facebook to WhatsApp and WhatsApp to Facebook were displayed in

Table 4: How the usage of RecApps influenced the presence of the apps involved in the most common app switching behaviors of the participants. A series of Wilcoxon Signed-Ranks tests suggest that RecApps encouraged the repetition of these behaviors.

	Baseline		RecApps		p
	M	SD	M	SD	
{Instagram, WhatsApp} [#]	11.12	8.83	16.10	15.04	0.008
{Instagram, Facebook} [#]	3.81	3.35	6.49	8.51	0.010
{Facebook, WhatsApp} [#]	3.91	3.58	6.58	9.73	0.049

3,814 widgets.

We then analyzed whether the displayed widgets effectively promoted smart-
 575 phone usage sessions involving these applications. As reported in Table 4, we
 found that smartphone usage sessions that included Instagram and WhatsApp
 were significantly ($p < 0.05$) more common during the *RecApps* phase than in
 the *baseline* phase ($M = 16.10$, $SD = 15.04$ vs. $M = 11.12$, $SD = 8.83$, respec-
 tively). Significant differences ($p < 0.05$) also characterized smartphone usage
 580 sessions that included Instagram and Facebook, and Facebook and WhatsApp,
 respectively. This suggests that, while RecApps did not change the overall be-
 haviors of the participants with their smartphones, it encouraged the repetition
 of transitions between participants' favorite apps.

In many cases, in particular, the suggestions made by RecApps matched the
 585 users' most used apps. The recommendation algorithm adopted by our tool,
 however, accounts for other information as well, e.g., contextual factors like
 locations, performed activities, and day of the week. Therefore, it is ideally
 able to extract a) which transitions between the users' most used apps are more
 common, and b) when these transitions are more likely to happen. To fur-
 590 ther explore the benefits of mixing information about app usage and contextual
 information, we analyzed how a widget implementing a simpler recommenda-
 tion algorithm, one that always suggests the users' most used app, may have
 performed in comparison with RecApps. To this end, we first calculated the

“top N” apps used by each participant during the *baseline* phase. Then, we
595 measured:

- how often participants in the study were simply selecting one of their “top N” apps when using the widget, and
- how often participants in the study closed the widget to manually open one of their “top N” apps.

600 We found that, on average, each participant used the widget to select one of her “top N” app in 62.71 % of cases ($SD = 23.89\%$). However, in 32.29% of cases on average ($SD = 23.89\%$), participants used the widget to select an app that was not part of their most used apps. This means that RecApps is able to capture connections between rarely used apps that might be impossible
605 to detect with simpler algorithms, e.g., those suggesting users’ most used apps, only. On the contrary, by analyzing all the cases in which the visualized widget was not used, we found that participants manually opened a “top N” app that was not recommended in 36.33% of cases. Besides highlighting the need of further improving the accuracy of the recommendation algorithm adopted by
610 RecApps, this may suggest that there are cases in which the most used apps prevail over specific app switching behaviors. As reported by the participants in the final interview (see Section 4.3.4), a possible solution to take advantage of the ability of RecApps in detecting specific app switching behaviors while giving the necessary importance to the most used apps is to allow users to manually
615 specify a set of shortcuts that are always visualized by the widget.

4.3.4. *Advantages and Disadvantages*

By analyzing the data collected during the final interview, we found that only 4 participants had a neutral opinion about the usefulness of RecApps. One of them claimed to prefer “*existing gestures*” to switch between apps (P5), while
620 another found RecApps not useful since she typically did not use more than one app in a row. The majority of the participants, instead, found RecApps useful to reduce the time they needed to switch between their favorite apps. For P9, for

instance, RecApps “*speeded up the transitions between my apps,*” while P14 said that widgets were useful because they “*directly connected me with other apps.*”
625 Such a feature positively impacted the way participants used their smartphones, according to them:

“The app improved my smartphone usage. Before using RecApps, I had to exit from the current app and manually look for the next app to be opened. Now I spend less time to open an app.” (P2)

630 Interestingly, some participants asserted that they also used RecApps to “monitor” their smartphone usage:

“Besides reducing the time you need to switch between apps, RecApps allowed me to identify which apps I use more frequently, and it made me realize how I spend my time on the smartphone” (P15)

635 Knowing what your habitual switching behaviors are, in particular, can either encourage or dissuade users to click on a recommended item. On the one hand, indeed, RecApps helped participants to prioritize “*the apps I can’t do without*” (P1), without wasting their time on apps they rarely used. On the other hand, some suggestions reflected a habitual behavior that participants
640 would rather avoid:

“When I used Instagram, RecApps regularly suggested me to open Facebook. Sometimes this made me realize that I waste too much time on social media, so in that cases I typically forced myself not to click on the recommendation.” (P16)

645 This suggests the need of allowing users to further customize widget contents, e.g., to permanently disable the suggestion of some unwanted apps. As highlighted by P9, indeed, being able to customize the received recommendations, e.g., “*to manually define some useful shortcuts*”, would improve the overall experience with the app.

650 The obtrusiveness of the visualized widgets was the main concern of our participants. Two users described the visualized icons as too big for their screen,

while the fact that widgets were always present, even if closed, was described as “a little bit invasive” (P14), also because they increased the chances of “unintentional clicks” (P9). Participants, however, felt this problem as easily solvable, since RecApps obtrusiveness characterized only some specific moments of their smartphone usage, e.g., “while I was actively using another app” (P14) or “when I was using the keyboard” (P17). Accordingly, participants immediately suggested different improvements to overcome the identified issue. P12, for instance, would add the possibility of moving the widget in any part of the screen, not just vertically, while P7 would remove the widgets while the user is watching videos or photos. P2, instead, proposed a more complex solution: “you should find a way to automatically change the position of the widget depending on what the user is doing.”

5. Discussion

Despite a large body of HCI literature is dedicated to study smartphone usage under different aspects, Turner et al. [10] recently suggested the need of further exploring how to support “common application switching behaviour on smartphones.” While researchers have already demonstrated that smartphone users are likely to switch between different applications in the same usage session [4, 11, 12], little is known about how to support this kind of interaction on modern smartphones. In this paper we have shown evidence that habitual app switching behaviors can be automatically detected and proactively supported by means of the RecApps application. In this section, we first discuss our research questions by linking the retrieved findings to the existing literature on smartphone usage and interaction mechanisms. Then, we discuss the design space for proactively supporting app switching behaviors by referring to different design dimensions, ranging from visibility to cognitive and physical demand.

5.1. Detecting and Supporting App Switching Behaviors

The first step to support app switching interactions is to being able to detect the specific switches that are typically performed by the user (**RQ1**), in order

to facilitate their reproduction. In this respect, previous works only highlighted high-level trends characterizing these behaviors. Böhmer et al. [4], for instance, demonstrated that the first app that users open in a session typically belongs to the “communication” category, and that there exists a strong connection
685 between apps belonging to the “lifestyle” and “shopping” categories, respectively [4]. In this work, we demonstrated how the specific and *habitual* app switching behaviors of a user can be automatically extracted from her smartphone usage data through a data analytic methodology based on association rules mining (Section 3). Such an approach is able to mix different data, including
690 contextual information, to extract recurrent transitions between apps that are performed by the user under stable contextual situations. Differently from other approaches adopted in previous studies, where data of different users are mixed together to extract and analyze (typically off-line) global usage patterns, our approach is tailored for each user, and it can be applied in real-time on a
695 normal smartphone. The application of the approach on a smartphone usage dataset demonstrated the effectiveness of the approach, and it highlighted several types of app switching behaviors habitually performed by the users under different contexts, thus further motivating the need of exploring novel interactions to support these behaviors.

700 We therefore implemented the proposed approach in RecApps, an Android application designed as an interactive floating widget that proactively supports the transition from one app to another (**RQ2**, Section 4). The user study of RecApps highlighted both positive and negative aspects in the first version of the floating widget: while it supported the quick reproduction of habitual app
705 switching behaviors, especially between social networks and messaging apps, the cognitive demand of having to deal with a widget every time participants opened an app was sometimes a problem, especially when the visualized icons interfered with the current participants’ task. To our knowledge, RecApps is the first attempt to *proactively* support app switching behaviors. Several different
710 algorithms for app prediction have been proposed in the last years (e.g., [21, 32, 34]). On the one hand, however, most of these algorithms have mainly

been used for tasks that do not directly influence the user’s interaction, e.g., saving memory and execution time by pre-loading the right app [21, 13, 32]. On the other hand, researchers and practitioners have only explored very simple
715 approaches to exploit app-prediction algorithms and support users in opening the next app, e.g., by statically showing suggested apps in home-screen menus (see the work of Sun et al. [33], and also the recommendations provided by some of the today’s mobile user interfaces). As a consequence, users are often forced to explicitly stop the use of the current app, e.g., through a double tap on the home
720 button, and look for a new one to be opened, e.g., by browsing the list of apps that are currently opened. Similarly to the work of Banovic et al. [6], RecApps aims at supporting a proactive user’s interaction: app recommendations, in our case, dynamically appear *while* the user is interacting with her smartphone.

5.2. A Design Space for Supporting App Switching Behaviors

725 The RecApps user study and its findings allow us to explore several trade-offs that characterize the design space of supporting app switching behaviors in modern smartphones. To delimit the design space, we first present two new RecApps designs that take into account the suggestions of our participants (Figure 11(b) and Figure 11(c)). As requested in the user study, both the alternatives allow
730 users to customize recommendations, with the possibility to “lock” a suggested item, e.g., to have it always available when using a given app, or to disable it by clicking on the “x” icon the first time it appears, e.g., to avoid the recommendation of unwanted habitual switching behaviors.

735 The first alternative to the initial version, i.e., the floating design (Figure 11(b)), combines the benefits of having an immediate shortcut to the next apps always available with a floating and automatically adaptable position of the visualized icons. Recommended icons, in particular, can be separately dragged and dropped in any part of the smartphone screen. Furthermore, their position is
740 automatically adapted to the current user’s tasks. In Figure 11(b), for instance, suggestions are automatically moved above the keyboard when the user starts

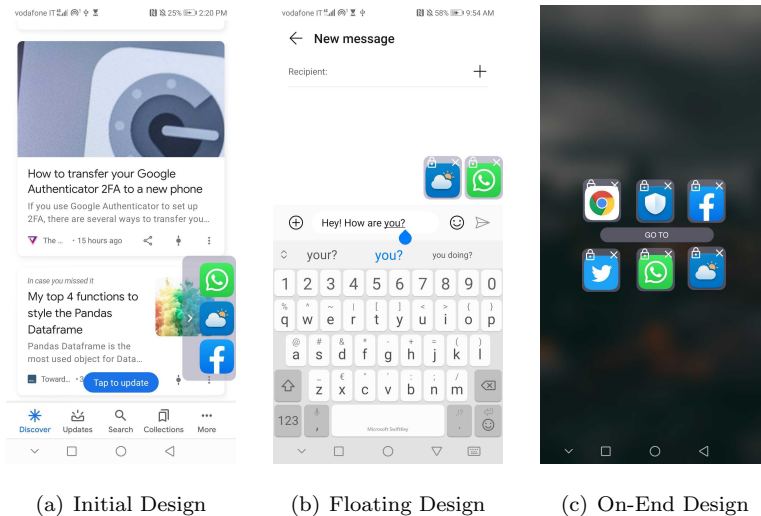


Figure 11: The initial design of RecApps (a) and two new RecApps designs (b and c).

writing a message, in such a way that the icons do not disturb neither the user’s writing nor the reading of the exchanged messages. Following the results of our user study, the design limits the number of visualized icons to two.

745

The second alternative (Figure 11(c)) is the on-end design. In this design, recommendations to switch to another app are presented through a dedicated window at the end of the interaction with the current app, only. The user can therefore invoke RecApps with any click or gesture that ends the usage of an app, e.g., through a click on the home button. This design emphasizes the need of limiting interruptions to the current task of the user, and it allows to maximize the number of recommendations the user can see at one time.

750

The three RecApps designs can be used to illustrate different trade-offs in our design space. We discuss them by referring to the design dimensions that Banovic et al. [6] adopted to explore how to proactively suggest app-specific tasks at the beginning of a smartphone session. Here, we exclude security and privacy issues since RecApps works “in the middle” of a smartphone session, and it does not

755

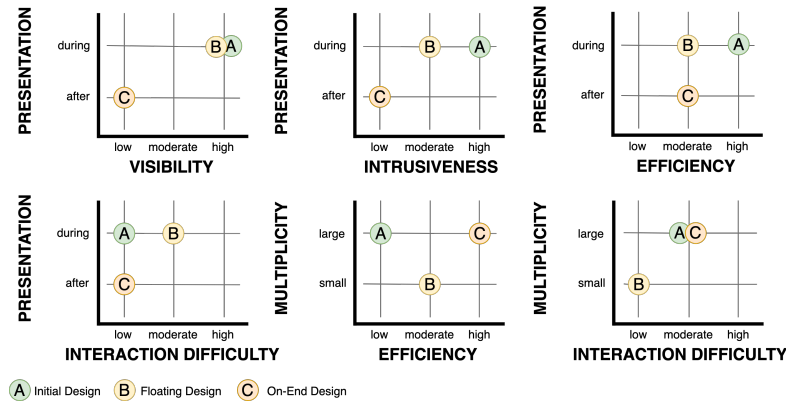


Figure 12: The three different RecApps designs mapped to the dimensions of the design space.

influence the functionality of secured lock screens. Figure 12 visually summarize
 760 the identified trade-offs, while each design dimension is discussed below.

Presentation strategy defines how and when to present recommended apps
 to the user. Our initial design and its floating version present always-on
 widgets that are visualized *during* the usage of other applications. The
 on-end design, instead, visualizes recommendations to switch to another
 765 app *after* the end of the interaction with the current app, e.g., when the
 user clicks on the home button. As reported in Figure 12, the chosen pre-
 sentation influences other design dimensions, from visibility to cognitive
 load.

Visibility refers to the level of engagement needed to reach a RecApps widget.
 770 The always-on widgets of the initial and floating designs are highly visible
 and provide an immediate shortcut towards other apps. Furthermore,
 they allow users to easily reproduce “back and forth” use cases, with
 users that can easily mix the usage of an app with secondary tasks, e.g.,
 to save a quick note about what they are reading on a web browser page.
 775 Recommended apps in the on-end design, instead, are not immediately
 visible, but they appear when the user is leaving the current app, only.

Intrusiveness express how much RecApps impact the time needed by users

to get to a recommended item. Intrusive designs can help increase visibility. However, as reported by our participants, the high intrusiveness of our initial design negatively impacted the usage of other apps by causing frequent unintentional touches. Furthermore, recommended icons interfered with some users' activities, e.g., writing on a keyboard, thus forcing users to close the widget or moving it away. The floating design aims at limiting these drawbacks by adapting the position of the recommended icons to the current user activity. Conversely, the on-end design minimizes intrusiveness, and it does not interfere with the usage of the current app.

Efficiency defines how quickly the user can switch to the next app from the current one. Participants of our user study appreciated the initial design of RecApps since it allowed them to rapidly switch between their favorite apps. A similar efficiency can be reached through the floating design. However, users could find it more difficult to reach a desired icon given the variable widget position. Similarly, the efficiency of the on-end design is influenced by the adopted presentation strategy, which forces users to perform an action, e.g., clicking the home button, before being able to switch to another app.

Interaction difficulty represents the difficulty of RecApps interaction in terms of cognitive and physical demand. In order to reduce the interaction difficulty, RecApps should present a limited number of icons in a easily accessible position of the screen, e.g., the initial bottom-right location of our initial design. However, such a position may not always be available, e.g., when writing a message. The floating design limits such an issue by automatically moving icons on the screen. However, users may find themselves confused by the continuous change of the icons' position, thus incurring in a higher cognitive demand. On the contrary, given the need of physically closing the current app, users may incur in a higher physical demand in the on-end design.

Multiplicity refers to how many different suggested apps are presented at any

given time. As reported in the last two charts of Figure 12, although a high number of icons might increase efficiency, e.g., with a higher probability of finding the desired app to switch to, it might also increase cognitive demand. Our user study, moreover, indicated that the initial design was likely to promote unintentional clicks when visualizing too many icons. That is the reason why, when using always-on widgets, RecApps should present a limited number of recommendations, e.g., as in the floating design of Figure 11(b). The number of suggestions could be increased when the visualized icons do not impact the usage of the current app, e.g., as in the on-end design.

Although we explored only a portion of the design space, researchers can take advantage of the described trade-offs to design future interactions supporting app switching behaviors that prioritizes different design dimensions.

6. Limitations and Future Works

There are some limitations to be considered in our work. Such limitations could inform future work in the design of novel interfaces to support app switching behaviors. Both our formative and user studies mainly involved students and young adults coming from the same geographical area and with a similar cultural background. We have to acknowledge that results may vary for different cultural settings and different ages: prior work such as Zhao et al. [39] has indeed shown that there is a very diverse set of smartphone users with different app usage patterns. Nevertheless, although the reported findings may not generalize, our work provides rich, qualitative design insights that we expect to be transferable [44] to the specific usage patterns of a given population.

Furthermore, RecApps supported app switching behaviors through app-level shortcuts, only, and the data analyzed and collected in this work do not allow to discriminate between app switching behaviors that are caused by usage routines and those that are due to interdependencies between apps. Multi-app sessions, however, often involve data transfer between apps, e.g., copying an address

from a chat window to a navigation app, or saving an image from Twitter and edit it in an image editor app. In this respect, many different apps on today’s smartphones offer a “share via” feature that allows users to integrate or transfer
840 content from one app into another. On the one hand, this implies that, during the study, RecApps may have been redundant for certain type of app switching. On the other hand, the effectiveness of integrating such “intent-level” shortcuts in RecApps, e.g., by taking into account other app-based features in the recommendation algorithm, could be explored in future works. An example of a
845 tool that integrates “intent-level” shortcuts is the *MessageOnTap* [45] system, that uses the text extracted from a chat conversation to suggest task shortcuts, e.g., seeing photos taken at Los Angeles, that can likely streamline next actions. Ideally, transferring data while going directly to the right page in the target app would make tools like RecApps even more effective. Besides investigating
850 intent-level shortcuts, other recommendations algorithms could be also explored, e.g., to increase serendipity while assisting users in repeating their habitual app switching behaviors. Furthermore, our design space could be further exploited to quantitatively investigate, e.g., through user evaluations, which are the most effective modalities to present and visualize recommendations. Finally, we also
855 see promise in exploring how to suggest *meaningful* app switching behaviors, only. As highlighted by some of our participants, indeed, RecApps sometimes suggested habitual app switching behaviors that they would rather suppress: a new version of RecApps could therefore promote the development of meaningful mobile interactions and, consequently, discouraging habitual patterns that
860 negatively influence users’ digital wellbeing.

7. Conclusions

App switching behaviors, i.e., transitioning between different apps in a smartphone usage session to consume content, are common among users. Unfortunately, today’s smartphones offer limited support to transition from one app to
865 another. In this paper, we have investigated how these behaviors can be au-

tomatically detected, and how modern smartphones could proactively support their reproduction. We proposed an approach based on association rules mining to extract recurrent links between contextual cues and used apps from smartphone usage data. Furthermore, we proposed different design opportunities to improve mobile device use by supporting app switching behaviors through the proactive suggestion of the next apps to be used. Our RecApps prototype, in particular, analyzes previous users' interaction with the smartphone, and it is able to suggest habitual app transitions for different contexts through a floating widget. Our work provides motivation for the design of novel interfaces to support app switching behaviors, and opens up many opportunities for future work.

References

- [1] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, D. Estrin, Diversity in smartphone usage, in: Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 179–194. doi:10.1145/1814433.1814453.
- [2] A. Oulasvirta, T. Rattenbury, L. Ma, E. Raita, Habits make smartphone use more pervasive, *Personal and Ubiquitous Computing* 16 (1) (2012) 105–114. doi:10.1007/s00779-011-0412-2.
- [3] L. Barkhuus, V. Polichar, Empowerment through seamfulness: Smart phones in everyday life, *Personal and Ubiquitous Computing* 15 (2011) 629–639. doi:10.1007/s00779-010-0342-4.
- [4] M. Böhmer, B. Hecht, J. Schöning, A. Krüger, G. Bauer, Falling asleep with angry birds, facebook and kindle: A large scale study on mobile application usage, in: Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services, Mobile-

HCI '11, Association for Computing Machinery, New York, NY, USA, 2011, p. 47–56. doi:10.1145/2037373.2037383.

- 895 [5] D. Hintze, P. Hintze, R. D. Findling, R. Mayrhofer, A large-scale, long-term analysis of mobile device usage characteristics, *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1 (2) (2017). doi:10.1145/3090078.
- [6] N. Banovic, C. Brant, J. Mankoff, A. Dey, Proactivetasks: The short of mobile device use sessions, in: *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services, MobileHCI '14*, Association for Computing Machinery, New York, NY, USA, 2014, p. 243–252. doi:10.1145/2628363.2628380.
- 900 [7] V. Srinivasan, S. Moghaddam, A. Mukherji, K. K. Rachuri, C. Xu, E. M. Tapia, Mobileminer: Mining your frequent patterns on your phone, in: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '14*, Association for Computing Machinery, New York, NY, USA, 2014, p. 389–400. doi:10.1145/2632048.2632052.
- 905 URL <https://doi.org/10.1145/2632048.2632052>
- [8] K. Lukoff, C. Yu, J. Kientz, A. Hiniker, What makes smartphone use meaningful or meaningless?, *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2 (1) (2018) 22:1–22:26. doi:10.1145/3191754.
- 915 [9] D. Hintze, R. D. Findling, M. Muaaz, S. Scholz, R. Mayrhofer, Diversity in locked and unlocked mobile device usage, in: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, UbiComp '14 Adjunct*, Association for Computing Machinery, New York, NY, USA, 2014, p. 379–384.
- 920 doi:10.1145/2638728.2641697.
- URL <https://doi.org/10.1145/2638728.2641697>

- [10] L. D. Turner, R. M. Whitaker, S. M. Allen, D. E. J. Linden, K. Tu, J. Li, D. Towsley, Evidence to support common application switching behaviour on smartphones, *Royal Society Open Science* 6 (2019). doi:10.1098/rsos.190018.
- 925
- [11] S. L. Jones, D. Ferreira, S. Hosio, J. Goncalves, V. Kostakos, Revisitation analysis of smartphone app use, in: *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 1197–1208. doi:10.1145/2750858.2807542.
- 930
- [12] J. P. Carrascal, K. Church, An in-situ study of mobile app & mobile search interactions, in: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 2739–2748. doi:10.1145/2702123.2702486.
- 935
- [13] K. Huang, C. Zhang, X. Ma, G. Chen, Predicting mobile application usage using contextual information, in: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing, UbiComp '12*, Association for Computing Machinery, New York, NY, USA, 2012, p. 1059–1065. doi:10.1145/2370216.2370442.
- 940
- [14] P. Leroux, K. Roobroeck, B. Dhoedt, P. Demeester, F. De Turck, Mobile application usage prediction through context-based learning, *Journal of Ambient Intelligence and Smart Environments* 5 (2) (2013) 213–235.
- [15] R. Baeza-Yates, D. Jiang, F. Silvestri, B. Harrison, Predicting the next app that you are going to use, in: *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 285–294. doi:10.1145/2684822.2685302.
- 945
- [16] A. Monge Roffarello, L. De Russis, Towards detecting and mitigating smartphone habits, in: *Adjunct Proceedings of the 2019 ACM International*
- 950

Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers, UbiComp/ISWC '19 Adjunct, Association for Computing Machinery, New York, NY, USA, 2019, p. 149–152. doi:10.1145/3341162.3343770.

955 [17] How do i disable suggested apps on my samsung device?, accessed: 2020-07-17 (2020).

URL <https://www.samsung.com/au/support/mobile-devices/how-do-i-disable-suggested-apps-on-samsung-android-pie/>

[18] T. Sohn, K. A. Li, W. G. Griswold, J. D. Hollan, A diary study of mobile
960 information needs, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08, Association for Computing Machinery, New York, NY, USA, 2008, p. 433–442. doi:10.1145/1357054.1357125.

URL <https://doi.org/10.1145/1357054.1357125>

965 [19] H. Verkasalo, Contextual patterns in mobile service usage, Personal and Ubiquitous Computing 13 (2009) 331–342. doi:10.1007/s00779-008-0197-0.

[20] T. Do, J. Blom, D. Gatica-Perez, Smartphone usage in the wild: A large-scale analysis of applications and context, 2011, pp. 353–360. doi:10.
970 1145/2070481.2070550.

[21] T. Yan, D. Chu, D. Ganesan, A. Kansal, J. Liu, Fast app launching for mobile devices using predictive user context, in: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, MobiSys '12, Association for Computing Machinery, New York, NY, USA,
975 2012, p. 113–126. doi:10.1145/2307636.2307648.

URL <https://doi.org/10.1145/2307636.2307648>

[22] A. Oulasvirta, S. Tamminen, V. Roto, J. Kuorelahti, Interaction in 4-second bursts: The fragmented nature of attentional resources in mobile hci, in:

- 980 Proceedings of the SIGCHI Conference on Human Factors in Computing
Systems, CHI '05, Association for Computing Machinery, New York, NY,
USA, 2005, p. 919–928. doi:10.1145/1054972.1055101.
- [23] A. Bianchi, J. Phillips, Psychological predictors of problem mobile phone
use, *Cyberpsychology & behavior : the impact of the Internet, multimedia
and virtual reality on behavior and society* 8 (2005) 39–51. doi:10.1089/
985 cpb.2005.8.39.
- [24] U. Lee, J. Lee, M. Ko, C. Lee, Y. Kim, S. Yang, K. Yatani, G. Gweon,
K.-M. Chung, J. Song, Hooked on smartphones: An exploratory study
on smartphone overuse among college students, in: *Proceedings of the
32Nd Annual ACM Conference on Human Factors in Computing Sys-
990 tems, CHI '14, ACM, New York, NY, USA, 2014, pp. 2327–2336. doi:
10.1145/2556288.2557366.*
- [25] S. Lanette, P. K. Chua, G. Hayes, M. Mazmanian, How much is 'too
much'?: The role of a smartphone addiction narrative in individuals' ex-
perience of use, *Proceedings of the ACM on Human-Computer Interaction*
995 *2 (CSCW)* (2018) 101:1–101:22. doi:10.1145/3274370.
- [26] R. Larose, J.-H. Kim, W. Peng, *Social Networking: Addictive, Compulsive,
Problematic or Just Another Media Habit?*, 2011, pp. 59–81.
- [27] J. A. Tran, K. S. Yang, K. Davis, A. Hiniker, Modeling the engagement-
disengagement cycle of compulsive phone use, in: *Proceedings of the 2019
1000 CHI Conference on Human Factors in Computing Systems, CHI '19, ACM,
New York, NY, USA, 2019, pp. 312:1–312:14. doi:10.1145/3290605.
3300542.*
- [28] A. J. van Deursen, C. L. Bolle, S. M. Hegner, P. A. Kommers, Modeling
habitual and addictive smartphone behavior: The role of smartphone usage
1005 types, emotional intelligence, social stress, self-regulation, age, and gender,
Computers in Human Behavior 45 (2015) 411 – 420. doi:https://doi.
org/10.1016/j.chb.2014.12.039.

- [29] M. Ko, S. Choi, K. Yatani, U. Lee, Lock n' lol: Group-based limiting assistance app to mitigate smartphone distractions in group activities, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16, ACM, New York, NY, USA, 2016, pp. 998–1010. doi:10.1145/2858036.2858568.
- [30] C. Shin, A. K. Dey, Automatically detecting problematic use of smartphones, in: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, ACM, New York, NY, USA, 2013, pp. 335–344. doi:10.1145/2493432.2493443.
- [31] C. Tossell, P. Kortum, A. Rahmati, C. Shepard, L. Zhong, Characterizing web use on smartphones, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, ACM, New York, NY, USA, 2012, pp. 2769–2778. doi:10.1145/2207676.2208676.
- [32] A. Parate, M. Böhmer, D. Chu, D. Ganesan, B. M. Marlin, Practical prediction and prefetch for faster access to applications on mobile phones, in: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 275–284. doi:10.1145/2493432.2493490.
- [33] C. Sun, J. Zheng, H. Yao, Y. Wang, D. F. Hsu, Apprush: Using dynamic shortcuts to facilitate application launching on mobile devices, *Procedia Computer Science* 19 (2013) 445 – 452, the 4th International Conference on Ambient Systems, Networks and Technologies (ANT 2013), the 3rd International Conference on Sustainable Energy Information Technology (SEIT-2013). doi:<https://doi.org/10.1016/j.procs.2013.06.060>.
- [34] X. Zou, W. Zhang, S. Li, G. Pan, Prophet: What app you wish to use next, in: Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication, UbiComp '13 Adjunct, Association for Computing Machinery, New York, NY, USA, 2013, p. 167–170. doi:10.

1145/2494091.2494146.

URL <https://doi.org/10.1145/2494091.2494146>

- [35] N. Natarajan, D. Shin, I. S. Dhillon, Which app will you use next?
1040 collaborative filtering with interactional context, in: Proceedings of the
7th ACM Conference on Recommender Systems, RecSys '13, Association
for Computing Machinery, New York, NY, USA, 2013, p. 201–208.
doi:10.1145/2507157.2507186.
- [36] Firebase helps mobile and web app teams succeed, accessed: 2020-07-17
1045 (2020).
URL <https://firebase.google.com/>
- [37] R. Agrawal, R. Srikant, Fast algorithms for mining association rules in
large databases, in: Proceedings of the 20th International Conference on
Very Large Data Bases, VLDB '94, Morgan Kaufmann Publishers Inc., San
1050 Francisco, CA, USA, 1994, pp. 487–499.
- [38] P.-N. Tan, M. Steinbach, V. Kumar, Introduction to Data Mining, (First
Edition), Addison-Wesley Longman Publishing Co., Inc., Boston, MA,
USA, 2005.
- [39] S. Zhao, J. Ramos, J. Tao, Z. Jiang, S. Li, Z. Wu, G. Pan, A. K. Dey,
1055 Discovering different kinds of smartphone users through their application
usage behaviors, in: Proceedings of the 2016 ACM International Joint Con-
ference on Pervasive and Ubiquitous Computing, UbiComp '16, Associa-
tion for Computing Machinery, New York, NY, USA, 2016, p. 498–509.
doi:10.1145/2971648.2971696.
- [40] Adapt your app by understanding what users are doing, accessed: 2020-
1060 07-17 (2020).
URL [https://developers.google.com/location-context/
activity-recognition/](https://developers.google.com/location-context/activity-recognition/)

- [41] Coronavirus disease (covid-19) pandemic, accessed: 2020-07-17 (2020).
1065 URL [https://www.who.int/emergencies/diseases/
novel-coronavirus-2019](https://www.who.int/emergencies/diseases/novel-coronavirus-2019)
- [42] Zoom meetings & chat, accessed: 2020-07-17 (2020).
URL <https://zoom.us/>
- [43] Cloud firestore: Store and sync app data at global scale, accessed: 2020-
1070 07-17 (2020).
URL <https://firebase.google.com/products/firestore>
- [44] L. Krefting, Rigor in qualitative research: the assessment of trustworthi-
ness., *The American journal of occupational therapy : official publication
of the American Occupational Therapy Association* 45 3 (1991) 214–22.
1075 doi:10.5014/ajot.45.3.214.
- [45] F. Chen, K. Xia, K. Dhabalia, J. I. Hong, Messageontap: A suggestive
interface to facilitate messaging-related tasks, in: *Proceedings of the 2019
CHI Conference on Human Factors in Computing Systems, CHI '19, As-
sociation for Computing Machinery, New York, NY, USA, 2019, p. 1–14.*
1080 doi:10.1145/3290605.3300805.