

# What and How long: Prediction of Mobile App Engagement

YUAN TIAN, University of Nottingham, United Kingdom

KE ZHOU, University of Nottingham, United Kingdom

DAN PELLEGG, Yahoo Research, Israel

User engagement is crucial to the long-term success of a mobile app. Several metrics, such as dwell time, have been used for measuring user engagement. However, how to effectively predict user engagement in the context of mobile apps is still an open research question. For example, do the mobile usage contexts (e.g., time of day) in which users access mobile apps impact their dwell time? Answers to such questions could help mobile operating system and publishers to optimize advertising and service placement. In this paper, we first conduct an empirical study for assessing how user characteristics, temporal features, and the short/long-term contexts contribute to gains in predicting users' app dwell time on the population level. The comprehensive analysis is conducted on large app usage logs collected through a mobile advertising company. The dataset covers more than 12K anonymous users and 1.3 million log events. Based on the analysis, we further investigate a novel mobile app engagement prediction problem – can we predict simultaneously what app the user will use next and how long he/she will stay on that app? We propose several strategies for this joint prediction problem and demonstrate that our model can improve the performance significantly when compared with the state-of-the-art baselines. Our work can help mobile system developers in designing a better and more engagement-aware mobile app user experience.

CCS Concepts: • **Information systems** → **Mobile information processing systems**; • **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**.

Additional Key Words and Phrases: mobile apps, user engagement, app usage, dwell time, next app prediction, app engagement prediction, demographics, behavior modeling, user modeling

## 1 INTRODUCTION

Mobile devices have become an increasingly ubiquitous part of our everyday life. People access mobile apps to fulfil their information needs or accomplish their daily tasks. Users are generally engaged with an app when they appreciate the mobile app content to which they have given their attention. Properly monitoring user engagement is one of the key ingredients of success to improve user experience and retention. One way user engagement has been measured at a large scale is by tracking how long users spend with content, e.g., the time spent on a webpage. Studies on user engagement in the contexts of desktop-based systems [56] and websites [13, 65] have shown that simple metrics such as dwell time are meaningful and robust in modelling user engagement. More importantly, these studies have shown that with an awareness of engagement, users' experience with a system can be substantially improved which in turn leads to user growth, user retention, and increasing revenue streams.

---

Authors' addresses: Yuan Tian, yuan.tian@nottingham.ac.uk, University of Nottingham, Wollaton Road, Nottingham, NG8 1BB, United Kingdom; Ke Zhou, ke.zhou@nottingham.ac.uk, University of Nottingham, Wollaton Road, Nottingham, NG8 1BB, United Kingdom; Dan Pelleg, pellegd@acm.org, Yahoo Research, Haifa, Israel.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1046-8188/2021/1-ART1 \$15.00

<https://doi.org/10.1145/3464301>




Fig. 1. Next app and app dwell time prediction.

However, even many past works in mobile computing have investigated how individuals download, install and use different apps on their mobile devices [26, 40, 49, 66], to our knowledge, few studies have examined how to effectively predict how long users would stay with a mobile app. Real-world mobile app usage behaviour is a complex phenomenon driven by a number of competing factors. Game apps, in general, have a higher probability to be used for a long period, whereas weather app is, not surprisingly, shorter. Intuitively, there are certain times in a day when a user might be more likely to engage with certain mobile apps: John might be more likely to engage with game apps for a longer time at night after work. User characteristics could also make a difference: female users may spend longer time with shopping apps than male users [44]. Despite that the averaged total time spent on different app categories were reported in [6, 32, 60], there is little research that comprehensively analyzes the dwell time of mobile apps. This motivates our first research question (**RQ1**):

- *What are the factors (user characteristics and contexts) that influence the dwell time a user spends on an app?*

As far as we know, there have been many works conducted for modelling users' behaviour on choosing one particular app under different contexts. However, they do not describe how users engage with that app. Attention is a scarce resource in the modern world. For instance, a user may become immersed in the video watching or quickly abandon it – the distinction of which will be clear if we know how much time the user spent interacting with this app given different contexts. Next app prediction is characterised as the willingness to use an app, whereas engagement is the usage pattern after accessing the app. Though most researchers have focused on measuring which app user will use [2, 34, 35, 61], the engagement of apps is still not well understood. Hence, we consider which app user will use and the engagement level of how long the user will stay with this app an aggregated measure of users' app usage behaviour. Furthermore, app engagement (dwell time) is much dependent on the app content itself. As we mentioned above, checking the weather app is always shorter than playing games. Therefore, it is meaningless to predict how long a user will stay regardless of which app user is engaging. Given the inter-dependency between an app and app dwell time, we aim to predict the next app as well as app engagement (dwell time) as shown in Figure 1. Therefore our second research question (**RQ2**) is:

- *Can we simultaneously predict which app user will use next and how long the user will stay on this app?*

Answers to these questions have a profound impact on the success of an app, as engagement awareness can radically improve users' experience with digital services [45].

To answer **RQ1**, we first demonstrate how different factors affect users' app dwell time by presenting the first population-level analysis of how different contexts affect the app usage duration. The analysis is based on a large-scale mobile app usage dataset with more than 1.3 million logs from over 9K unique apps and 12K users. The data was collected from a mobile advertising company over a period of one week. Specifically, we consider the influential factors from two aspects: user characteristics (e.g., age, gender, device type, historical preferences) and context (e.g., hour, weekday, last used app, periodic pattern). We find that, for example, users between 20 and 40 years old are more likely to have a shorter dwell time than teenagers and older people. Both the demographics and device type have an influence on how long user stay with mobile apps. Furthermore, we also establish that the app dwell time differs significantly across different time in a day. More importantly, we observe that users' app dwell time maintains periodic patterns and follows historical trends. For example, some users spend a similar length of time to regularly check the shopping apps every day (i.e., after every 24 hours). Additionally, users have different historical engagement habits on their app usage duration. For example, some users always prefer staying long on social apps, while others tend to only check for a short while.

Based on the comprehensive analysis of users' app dwell time, we are able to conduct the study of how users' app dwell time can be inferred from these features. We then set to answer **RQ2** – how to predict the next app and how long the user will stay on this app simultaneously? Based on past work on next app prediction [2, 34, 35], we propose several joint prediction models (sequential, stacking, and boosting) for such app engagement prediction, whereas the dwell time is represented as discrete levels (light, medium, and intensive) defined based on different app categories. Additionally, different from personalized models in prior work [2, 34, 35], our models leverage the community-behaviour patterns by extracting predictive features from the entire user population.

To summarize, our main contributions are two-folds:

- We conduct the first empirical analysis of mobile app engagement based on dwell time with a large-scale data set collected from thousands of users.
- Our research investigates a novel problem on simultaneously predicting which app user will use and how long the user will stay on that app. An effective boosting based prediction model is proposed, which outperforms 56.8% over the baseline approaches.

To the best of our knowledge, this is the first empirical study on inferring predictive features at the scale of millions of logs for app dwell time, assessing how user characteristics and context features impact the dwell time with a mobile app. Our proposed predictive models are empirically validated to be effective for this novel task.

The rest of the paper is organized as follows. We first review the prior literature in Section 2. Section 3 describes the dataset and presents a descriptive analysis of the app usage logs. Section 4 introduces all the features we extracted from our dataset and provides insights on how these features could impact app dwell time prediction. Section 5 presents our methodology for predicting the next app and app dwell time simultaneously, and we propose three joint prediction strategies to solve this problem. We then show the experimental results of our proposed models and analyze the effectiveness of different models in Section 6. We discuss the implications and potential limitations of our work in Section 7 and conclude in Section 8.

## 2 RELATED WORK

Our work is related to prior literature on user engagement analysis (§2.1) and next app/engagement prediction (§2.2).

### 2.1 User Engagements

**2.1.1 Engagement Measurements.** Approaches to measuring user engagement of online services can be divided into three main groups: (a) self-reported engagement, (b) cognitive engagement, and (c) online behaviour metrics [30]. In the first group (a), questionnaires and interviews are used to elicit user engagement attributes or to create user reports and to evaluate engagement [28]. The second type of approach (b) uses task-based methods and physiological measures to evaluate the cognitive engagement (e.g., facial expressions, vocal tone, and heart rate) using tools such as eye-tracking [16], heart rate monitoring, brain readings from headset [36], swipe on the screen [39] and mouse tracking [25]. However, these two types of methods have known drawbacks, e.g., reliance on user subjectivity of the self-reported engagement, and only be able to measure a small number of user interactions of the cognitive engagement.

The third type of approach (c), adopted by the web-analytics community, has been studying user engagement through online behaviour metrics that assess users' depth of engagement within a site, e.g., the time spent on a webpage. Studies on user engagement in the contexts of desktop-based systems [56] and websites [13, 65] have shown that simple metrics such as dwell time are meaningful and robust in modelling user engagement. For example, the time spent on a resource has been validated as an effective metric for measuring user engagement in the context of web search [1, 3], and recommendation tasks [64]. Kelly et al. [19, 27] consider dwelling times as an indicator of page relevance or user satisfaction during search engine interactions. Yi et al. [64] recommend designing dwell time based user engagement metrics and claim that this would enable them to extract better user engagement signals for training recommendation systems thereby optimizing for long term user satisfaction.

Therefore, we also propose to use the metric of time spent on mobile apps (dwell time) as our user engagement metrics within a large-scale dataset. For now, minimal research has been done for modelling mobile app dwell time from a large-scale dataset; only the basic aggregated statics on app usage time was reported. Falaki et al. [17] found that 90% of app usage sessions would be less than 6 minutes and Xu et al. [60] reported that the majority of total network access time for all apps is from 10 seconds to 1 hour for each subscriber in one week. Li et al. [32] reported usage time for different app categories in total and found that communication apps account for 49% cellular time against all apps. Böhmer [6] found that the Libraries & Demos apps (default Updater, Google Services Framework, etc.) have the longest average usage time from opening to closing. However, these summarized basic statistics of app usage time can not provide an in-depth understanding of what factors could influence the dwell time user spend on an app and whether we could predict how long user would stay with an app. In our work, we conduct the first empirical analysis of dwell time during app usage based on a large-scale data set collected from thousands of users; and we are the first to assess how different kinds of features (e.g., demographics, device type, hour of day, and last used app) impact the mobile app dwell time.

**2.1.2 Influential Factors of App Usage.** Lots of previous research has been conducted to uncover the factors that influence app usage behaviour, ranging from time, demographics, device types, last used app, to periodic patterns. Xu et al. [60] and Li et al. [32] discovered the temporal pattern of app usage, e.g., news apps were more frequently used in early mornings, whereas sports apps were more frequently used in evenings. They also found that, in general, the app usage frequency changed during the day, which grew from 6 am and reached its first peak around 11 am, and were

most active during the evening (7 pm to 9 pm). Additionally, demographics have been identified as an important factor that would impact the way users engage with their apps [29, 48, 53, 55, 68], e.g., how users select apps or make in-app purchases on their smartphones. Li and Lu [31] ever analyzed the device-specific apps usage patterns from large-scale android users and they found that users rely less on the cellular network as the price of device model increases. For the correlations between apps usage, Li et al. [6, 32, 60] all showed that the last used app would probably impact the next app the user is going to use. However, minimal research has been done for analyzing if any of these features could impact the dwell time of an app. In our work, we conduct the first empirical analysis of how long users will stay with an app on a large data set. We are the first to systematically model how user characteristics (e.g., demographics) and context features (e.g., time, last used app) impact the time user stay with an app.

## 2.2 App Usage and Engagement Prediction

**2.2.1 Next App and App Engagement Prediction.** Much research work has been conducted on predicting which app user will use (a.k.a. next app prediction). Based on the survey [9], we could know that before 2017, most of the app prediction/recommendation works are conducted during 2012-2013. Tan et al. [51] tried to treat users' app usage patterns as periodic time series cycles and predict the app will be used based on a prediction algorithm with fixed cycle length. Liao et al. [33] also predicted the next app based on mined temporal profiles for each app. Additionally, Huang et al [26], Shin et al. [49] and Zou et al. [70] all pointed out that the latest used app and time are more effective than location and other context information in the next app prediction. After 2017, neural approaches have become increasingly popular. Some researchers proposed different neural models for predicting the next app, including CNN [46] and LSTM [61], etc. Xu et al. [61] proposed a generic prediction model based on Long Short-term Memory (LSTM), to covert the temporal-sequence dependency and contextual information into a unified feature representation for next app prediction and stated that it outperforms other models. Additionally, Tian et al. [54] explored the prediction to identify if the pair of two app usage logs belong to the same task. As far as we know, no research has been done for predicting how long users will stay with an app. Only Mathur et al. [36] tried to model and predict users' attention-based engagements (two levels: focused attention and felt involvement) in the context of smartphones. They conducted the work based on physiological measures (headset readings) within 10 participants, and their research has no works regarding which app users are currently using. In our work, by leveraging a large-scale data set of users' app usage logs, we could be able to model users' engagement (dwell time) within specific apps by exploiting user characteristics, temporal context, and short/long-term behavioural patterns. Most importantly, we investigate the challenges of simultaneously predicting which app user will use and how long the user will stay with this app.

**2.2.2 Engagement (Dwell Time) Prediction in Other Areas.** For the novel prediction problem proposed in our work, although no research has been conducted on the app dwell time prediction, some researchers investigated the engagement (dwell time) modelling and prediction in other areas, e.g., session-based recommendation (SBR) systems [5, 58, 69], videos watching [59], news and non-news pages reading [23, 47], and media streaming [57].

The SBR tasks aim to predict the next click/buying/dwell time based on users' interactions in a session, e.g., buying one item after viewing several products (within a commerce site). They stated that dwell time should be used as a proxy to user satisfaction of the clicked result since clicked through or not is not enough to identify the satisfaction of the user. In their scenario, researchers aim to predict which link/product user would click among a list of recommended similar results (under specific search query). Additionally, they assumed that the longer user stays with the clicked

result, the more satisfied the user will be with the result. However, in our dwell time prediction problem, a user would use an app occasionally with no recommendation context, and the longer user stays with this app does not directly mean the user is satisfied with it or not. The dwell time may be affected by the specific app (e.g., a weather app or game app), or whether the user accessed it during commuting or at night while they have more leisure time. Therefore, predicting how long a user will stay with an app is to predict the usage pattern while using the app. It would allow the service provider to optimize the user experience along with its business goals, the apps can be tuned to be more exploratory or exploitative based on the expected length of the usage. The SBR tasks can be solved by item-to-item and matrix factorization methods [22, 38], Markov Decision Process (MDP) based technique [52]. Recently, deep learning methods, Recurrent Neural Networks (RNN) have emerged as powerful methods of modelling sequential data in SBR [5, 58, 69].

The engagement pattern of the other online services, like video watching [59], news/non-news page reading [23, 47], and media streaming [57] have more similar characteristics of apps usage, i.e. how long user would stay could be affected by the category of the content or the original length of the video or news. Additionally, the underlying motivation for engagement prediction of these services is also the same as our work. They consider popularity and engagement as different measures of online behaviour. Although popularity describes the human behaviour of choosing one particular item, it does not describe how users would engage with this item. i.e., popularity is characterized as the willingness to click a video, whereas engagement is the video watching pattern after clicking. By knowing how long user would stay, we could be able to provide users with more satisfying content/services that increase long term user engagement and as a side-benefit. It also allows the service providers to optimize the user experience along with its business goals. Wu et al. [59] conducted a large-scale measurement study of engagement on 5.3 million videos over a two-month period and measured a set of engagement metrics (e.g., watch time, watch percentage) for online videos. They predicted engagement from video context, topics, and channel reputation, etc. Seki et al. [47] and Homma et al. [23] all clarified the characteristics of relationships between dwell time on news/non-news pages reading in order to discover which features are effective for predicting the dwell time, including (1) Dwell time by Device: desktops and mobiles; (2) Dwell time by access time; (3) Dwell time by if users visited from inside or outside the site; (4) Dwell time by click and non-click: if the user clicked links in the page; (5) Dwell time by scroll depth. Vasiloudis et al. [57] explored the prediction of session length in a mobile-focused music streaming service. They predicted the length of a session using contextual and user-based features including gender, age, subscription status, device, network type, duration of the user's last session, and time elapsed since the last session.

Given the task similarity between these engagement prediction studies and our focus (app engagement prediction), we selected two of them [57, 59] with the most features that could be extracted from our dataset as the baselines for comparing the performance of our proposed model regarding the app engagement prediction. We all have similar engagement characteristics and the motivation for engagement prediction, e.g., the engagement (dwell time) is originally correlated with the content category and more exploratory or exploitative service could be provided based on the expected length of the usage.

In summary, for the traditional next app prediction task, we selected four models from previous works that could be fitted to our dataset as baselines, include three works [49, 51, 70] based on the temporal pattern and contextual features and one recent work based on neural approach (LSTM) [61]. For the joint prediction problem of predicting the next app and engagement level together, since there is no existing similar joint prediction work (app dwell time has a high dependency on which app user is engaging), we added those two recent works [57, 59] for predicting dwell time as baselines. The only difference is they solely predict how long a user will stay based on the

Table 1. Overall statistics of the dataset.

OS	%Logs	%Users	Device	%Logs	%Users
android	78.0%	57.5%	Phone	94.3%	90.0%
ios	22.0%	42.5%	Tablet	5.7%	10.0%
Age	%Logs	%Users	Gender	%Logs	%Users
13-17	5.7%	9.0%	female	44.1%	51.8%
18-24	13.9%	18.4%	male	55.9%	48.2%
25-34	32.2%	30.3%			
35-54	43.2%	36.5%			
55+	5.0%	5.8%			




Fig. 2. Visualisation of app usage logs and the corresponding app usage duration. Five minutes of inactivity is used for segmenting mobile sessions.

specific item, without predicting on which item the user will engage with. In order to make them comparable baselines to our work, we assumed the ground truth of the next app is known for those two prior works [57, 59] and leveraged them specifically for predicting engagement (dwell time) of the known next app. This demonstrates the upper bound of those approaches (oracle performance) regarding the joint prediction problem.

### 3 DATA AND DESCRIPTIVE ANALYSIS

We start by describing the data used in this study, followed by some fundamental analysis that demonstrates the characteristics of our data.

#### 3.1 Dataset

The dataset used in our paper is collected from a mobile advertising company, a library that mobile developers integrate into their apps to measure app usage and allow in-app advertising. We collected a sample of logs from a week in March 2017 of more than 1.3 million logs with over 9K different apps and 12K users from the United States. Each log consists of the user’s general app usage information, such as demographics, timestamp, app category, app id and time spent. Each app belongs to one of 45 categories ranging from social, communication to business, etc. The definition of these app categories is consistent with the Google Play App taxonomy [21]. Table 1 shows various statistics of our dataset. Either Android or iOS operates the devices. 51.8% of app users are female, and most logs (over 70%) are generated by users between 25 and 54 years old.<sup>1</sup>

Following previous work [10], we consider a five-minute range of inactivity as the signal of ending a session as shown in Figure 2. If the user leaves an app but revisits any app within 5 minutes the *session* continues; otherwise, the *session* ends. In our work, we aim to predict which app the user will use next and how long the user will stay with that app. The app usage duration we aim to predict is calculated as follows (as shown in Figure 2): we aggregate the consecutive app

<sup>1</sup>Users have been classified into five age ranges in our dataset: 13-17,18-24,25-34,35-54 and 55+.

Table 2. Top 10 popular app categories.

App Categories	%Sessions	App Function Examples
productivity	28.6%	mail, calendar, notepad
social	10.6%	SNS, dating
tools	8.3%	calculator, screen lock, light
communication	6.8%	SMS, IM, free video calls
entertainment	5.1%	TV player, streaming video
utilities	3.3%	network, cleaner
sports	3.1%	live sports, sports news
music	2.7%	music player
lifestyle	2.5%	diary, discount, recipe
arcade	2.4%	games




Fig. 3. Overall average app usage duration for different app categories.

usage duration of the same app within each session. The app usage duration can be also referred to as app *dwell time*. We adopt this definition as the unit of our analysis for the rest of the paper. Note that we only consider user-triggered events; i.e. we do not include events that are triggered by background refresh when conducting our analysis. To reduce bias from users with a low level of engagement, we restricted our sample to those users who interacted with apps from at least five different categories. All the data was anonymized by removing all personally identifiable data prior to processing.

### 3.2 Distributions of App Usage

We start by presenting the distributions of the popularity of apps. The popularity of an app category can be measured using the total number of sessions. Table 2 shows the most popular app categories in our dataset. Generally, users are more likely to interact with app categories like productivity,






Fig. 4. PDF and CDF of app usage duration across all apps.

social, tools, communication, and entertainment apps nowadays. We also illustrate several examples of the most popular functionalities of the corresponding apps within each category. Additionally, we show the average app usage duration for different app categories in Figure 3, which ranges from less than 1 minute to over 10 minutes. We find that users always spend longer time when they are engaging with some app categories for relaxation, such as comics and games apps (i.e., card, word, puzzle, and board). It also states that different app categories initially have different lengths of app usage duration, e.g., game apps have a much longer duration than communication apps. On average for all the app categories, our users’ app dwell time lasted about 2.5 minutes. In Figure 4, we show the probability density function (PDF) and cumulative distribution function (CDF) of app usage duration. We can find that most of the app usage (93.7%) is less than 10 minutes and 80% of them only last less than 2 minutes.

We then look at how users engage with their apps throughout the day. Figure 5 plots the scaled amount of sessions (using min-max normalization) and average app usage duration each hour. We can find that the total app usage (in terms of session amount) is at its maximum in the afternoon and evening, peaking at around 7 PM. This aligns with findings reported in [6, 55]. Figure 5 also shows the average app dwell time regarding the different hours of the day. Generally, the average app usage duration across the day is between 2 minutes and 3 minutes. The duration increases after 6 PM and drops after 1 AM. This might be due to people have more leisure time during the non-working hours. We indeed find that the app categories associated with longer duration at late night are game and tool (system cleaner/VPN) apps.

### 3.3 Engagement Level Definition

To evaluate the performance of app usage duration prediction more intuitively, we propose to represent the time length of each duration as a discrete value, i.e., we classify the app usage duration into three engagement levels: light, medium, and intensive. Additionally, as we have found in Figure 3 that different app categories initially have different lengths of app usage duration, so it may not be reasonable to label the engagement levels without differentiating app categories. Specifically, to assign the corresponding engagement level of each app usage duration: (1) we first calculate the quantiles (33% and 67%) of the duration for different app categories respectively; (2) then we assign the level label to each duration based on their corresponding app category quantiles. For example, the 33% and 67% quantiles of weather apps are 6 seconds and 23 seconds respectively.




Fig. 5. App usage patterns across a day.

Table 3. Top 5 app categories with the biggest gender effects – the higher the gender effect is, the bigger difference exists in the app usage duration between male and female users. The gender group  $G_1$  has longer app usage duration.

App Category	Gender Group $G_1$	Gender Effect	App Category	Gender Group $G_1$	Gender Effect
widgets	female	3.30	casino	male	1.81
medical	female	2.54	video	male	1.64
entertainment	female	2.10	health-and-fitness	male	1.60
communication	female	1.64	finance	male	1.30
shopping	female	1.63	navigation	male	1.27

If the current usage duration of the weather app is 10 seconds, then its engagement level is medium. For the game apps whose 33% and 67% quantiles are 1.3 minutes and 5.4 minutes respectively. When a game app usage duration is longer than 5.4 minutes, its engagement level is intensive.

## 4 INFERRING USERS' APP DWELL TIME

To model the app engagement accurately, we need to uncover what information could be indicative factors for users' app dwell time. This is the main focus of this section.

### 4.1 User and Device Characteristics

**4.1.1 User Demographics.** Prior studies [29, 48, 68] have shown that demographics (age and gender) have a significant impact on how users select apps or make in-app purchases on their smartphones. However, little is known about the correlation between users' demographic information and the time spent when using an app. Establishing such a relationship is the main focus of this subsection. By analyzing our user mobile app usage data, we demonstrate that on average, the app usage duration for female users is 2.8 minutes, which is longer than the male users who spent 2.2 minutes. To show the general pattern of usage duration across all apps, Figure 6a presents the CDF of app usage duration for both the male and female users. We find that the duration distributions are similar between male and female users.

To investigate the effect of age, we demonstrate in Figure 6b that age affects app usage duration more significantly. Users between 25 and 54 years spend less time while engaging with apps than

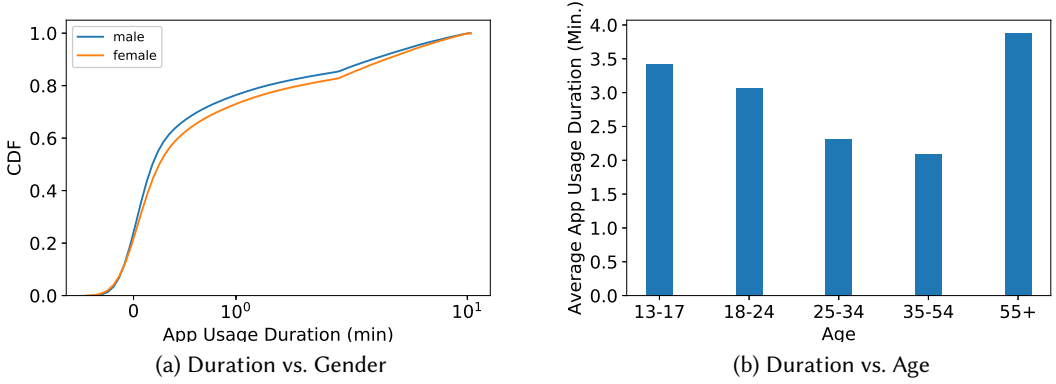


Fig. 6. Effect of gender and age on app usage duration.

those who are teenagers or seniors. This scenario is consistent with results from the previous user studies [18, 43]. Ferreira et al. [18] pointed out that micro-usage (less than 15 seconds) is popular across their participants aged between 22 and 40 years old. Pielot et al. [43] found that their participants (aged between 24 and 43 years old) need to deal with the large volume of notifications coming from personal communication due to high social expectations and the exchange of time-critical information.

Besides these findings of general patterns across all apps, we hypothesize that the impacts of demographics would vary across different app categories. Therefore, we further explore if users' app usage duration with each app category would be affected by their demographics. To measure the impacts brought by users with different genders in a more generalized manner, we calculate the *gender effect* for each app category of the two genders. First, for each app category  $C_a$ , we calculate the average app usage duration of female and male users respectively, i.e.,  $D(C_a, G_i)$  where  $G_i \in \{male, female\}$ . Secondly we calculate the *gender effect*,  $GE(C_a)$ , for each app category  $C_a$  as:

$$GE(C_a) = \frac{D(C_a, G_1)}{D(C_a, G_2)}$$

where  $GE(C_a)$  measures the *gender effect* of the app category  $C_a$ . The higher the effect is, the bigger difference exists in the app usage duration of the male and female users within this app category. The gender group in the numerator has a longer dwell time on this app category. Table 3 shows the top 5 app categories with the highest gender effects. We can find the female users have a longer app usage duration than male users with apps like widgets, medical, entertainment, communication, and shopping. On the other side, male users stay longer with the casino, video, health-and-fitness, finance, and navigation apps.

Similarly, we calculate the *age effect* brought by different users on the app dwell time. Since there are five age groups in our dataset, we choose the average app usage duration of all users as the reference value to calculate the *age effect*. Therefore, for each app category  $C_a$ , we calculate the average duration of users belonged to different age groups respectively, i.e.,  $D(C_a, A_i)$  where  $A_i \in \{13-17, 18-24, 25-34, 35-54, 55+\}$ . Then we calculate the *age effect*,  $AG(C_a)$ , for each app category  $C_a$  as:

$$AG(C_a) = \frac{D(C_a, A_i)}{D(C_a, A)}$$

Table 4. Top 5 app categories with the biggest age effects – the higher the age effect is, the bigger difference exists in the app usage duration between users with different age. The age group  $A_i$  has longer app usage duration when compared with other users.

App Category	Age Group $A_i$	Age Effects	App Category	Age Group $A_i$	Age Effects	App Category	Age Group $A_i$	Age Effects
shopping	13-17	2.75	widgets	18-24	2.35	food-and-drink	25-34	1.59
entertainment	13-17	1.81	entertainment	18-24	2.20	video	25-34	1.30
social	13-17	1.28	business	18-24	1.65	medical	25-34	1.30
photography	13-17	1.27	tools	18-24	1.41	music	25-34	1.29
sports	13-17	1.24	adventure	18-24	1.41	transportation	25-34	1.27
App Category	Age Group $A_i$	Age Effects	App Category	Age Group $A_i$	Age Effects			
lifestyle	35-54	1.29	productivity	55+	2.16			
casino	35-54	1.27	entertainment	55+	2.11			
word	35-54	1.23	board	55+	2.00			
travel	35-54	1.22	puzzle	55+	1.90			
action	35-54	1.21	books	55+	1.88			

which measures the *age effect* of the app category  $C_a$ .  $A$  represents the users across all ages who have engaged with the app category  $C_a$ . The higher the effect is, the bigger difference exists in the app usage duration of the users with corresponding ages  $A_i$ . The age group in the numerator has a longer dwell time on this app category when compared with other users. Table 4 shows the app categories with top *age effects*. It is interesting to find that teenage users have a longer duration in the shopping apps. Users between 25 and 34 prefer to stay longer with the food-and-drink apps than other users. It is not surprising to find that the older users over 55, will spend a much longer time when using the apps of entertainment, games (board/puzzle), and books apps since they have more spare time. For the longer time spent in productivity apps, this may result from that old users are not as efficient in the operations with the productivity apps (e.g., Microsoft Office, file managers).

**4.1.2 Device.** Li and Lu [31] ever analyzed the influence brought by different mobile device models on users' online time and they found that users rely less on the cellular network as the price of the device model increases. However, there has been no empirical research conducted for the impacts on dwell time with different apps brought by device characteristics. In this section, we compare app usage duration across different device types. There are mainly two different types of devices in our dataset: tablet and phone. We use the similar methodology of calculating the *gender effect* for each app category to calculate the *device effect* with the two device types. First, for each app category  $C_a$ , we calculate the average usage duration of phone users and tablet users respectively, i.e.,  $D(C_a, T_i)$  where  $T_i \in \{phone, tablet\}$ . Secondly, we calculate the *device effect*, for each app category  $C_a$  as:

$$DTE(C_a) = \frac{D(C_a, T_1)}{D(C_a, T_2)}$$

The higher the *device effect*  $DTE(C_a)$  is, the bigger difference exists in the app usage duration of the users with different device types. Table 5 shows the app categories with the most significant difference across the different devices. We can find that the tablet users have a longer dwell time on productivity, books, business, board, and travel apps. On the other hand, phone users have a longer dwell time with navigation, shopping, weather, family, and personalization apps.

## 4.2 Temporal Patterns

The temporal usage pattern is always an essential factor for inferring which app user will use, where we usually model the possibility of using different apps at a specific time [26, 33, 49]. From our descriptive analysis in Figure 5, we can find that the average app usage duration with all app

Table 5. Top 5 app categories with the biggest device effects – the higher the device effect is, the bigger difference exists in the app usage duration between the phone and tablet users. The users with device type  $T_1$  have longer app usage duration.

App Category	Device Type $T_1$	Device Effects	App Category	Device Type $T_1$	Device Effects
productivity	tablet	4.95	navigation	phone	8.04
books	tablet	3.29	shopping	phone	6.02
business	tablet	3.00	weather	phone	1.29
board	tablet	2.91	family	phone	1.22
travel	tablet	2.82	personalization	phone	1.19

Table 6. Index of dispersion of average app usage duration distribution across a day.

App Category	Index of Dispersion	App Category	Index of Dispersion
tools	0.02	comics	4.18
productivity	0.02	adventure (game)	3.86
utilities	0.02	strategy (game)	2.51
music	0.04	racing (game)	2.05
photography	0.05	family	1.59
news	0.06	transportation	1.35
widgets	0.07	medical	1.03




Fig. 7. Temporal patterns of average app usage duration for different app categories.

categories across a day is about 2 to 3 minutes, and no sharp fluctuations exist. We use the index of dispersion [12] as a normalized measure of the dispersion for the app usage duration distribution across a day, which is defined as the ratio of the variance  $\sigma^2$  to the mean  $\mu$ :  $D = \frac{\sigma^2}{\mu}$ . In general, the index of dispersion for all apps is 0.043. To validate that whether the temporal pattern varies on different app categories, we further explore the index of dispersion of the app usage duration across a day for different app categories respectively. Table 6 shows the app categories with the smallest and biggest index of dispersion in app usage duration. We can find that the usage duration of most game apps will be significantly different across a day. However, for some functional apps like productivity, tools, utilities, and photography, whenever the user accesses them, the app usage duration does not change much. Figure 7 illustrates the distribution of average usage duration across a day for several app categories, where we can find that besides the different variances, the

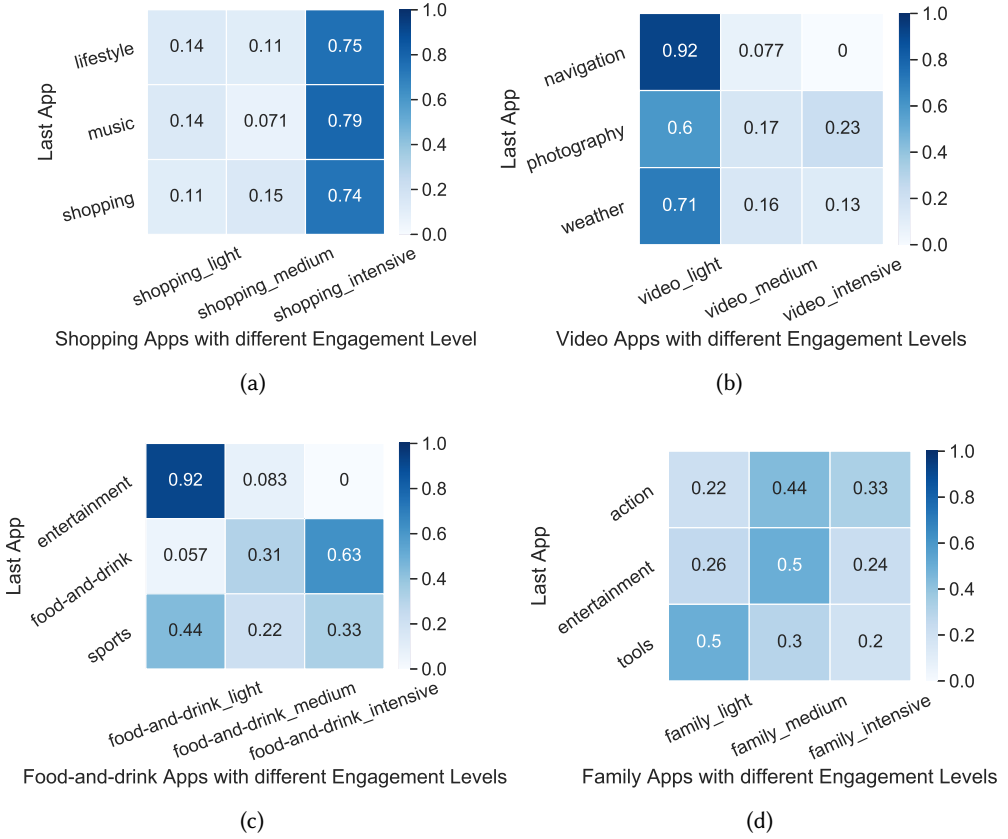


Fig. 8. Correlation between last used app and the next app's engagement level (discrete representation of app usage duration: light, medium and intensive). The darker color means higher probabilities that the next app will be engaged with the corresponding engagement level.

temporal patterns could be significantly different with various app categories (each app category has its own specific temporal pattern). For example, comics apps have a longer usage duration around 4 AM; the usage duration of racing apps is peaking around 3 AM; longer time is spent in adventure apps around 8 PM and 1 AM.

### 4.3 Short-term context

**4.3.1 Last Used App.** Previous studies identified that some apps were often used together [6, 32, 60]. For example, Li et al. [32] showed that some apps are installed together, whereas other studies [6, 60] found that some genres of apps are highly likely to be accessed sequentially. Therefore, we aim to explore whether the last used app could also impact the next app usage duration. To avoid the biases of the original differences of app usage duration across different app categories and measure the impacts more intuitively, we quantify the app usage duration into the corresponding engagement level defined as in Sec 3.3. Figure 8 shows several examples for illustrating the different patterns between specific last apps and the engagement levels of next apps. Each cell in the table represents the transition probability  $P_{ij}$  from the last used app  $a_i$  to the next app  $a_j$  at a given

engagement level  $e_{a_j}$ . By having  $P_{ij}$  for all the three engagement levels, we calculate its standard deviation as  $\sigma_j^i$ . This represents the extent of how much the last app  $a_i$  would have very different transition probabilities on different engagement levels of the next app  $a_j$ . The higher  $\sigma_j^i$ , the more likely the last app  $a_i$  would result in a specific engagement level of the next app  $a_j$ . For a given app  $a_j$ , we average  $\sigma_j^i$  across all the last used apps  $a_i$  as  $\sigma_j = \sum_{i=1..N} \sigma_j^i / N$ , where  $N$  is the amount of unique last used apps. For all the apps, the  $\sigma_j$  ranges between 0.11 and 0.32, whereas the median is 0.15. For the apps with higher  $\sigma_j$ , i.e., shopping apps ( $\sigma_j \approx 0.32$ ) and video apps ( $\sigma_j \approx 0.29$ ), the last used app could lead them to be used with specific engagement level (intensive for shopping and light for video). For example, for the video apps (as shown in Figure 8b), if they are used after using navigation apps, the time spent on the video app is much more likely to be short. This could be explained by users who are on their commute to work. After checking the navigation, users may need to wait for the bus or subways, so they could have time for enjoying a short video. Additionally, we also show the app categories with average  $\sigma_j$  and lower  $\sigma_j$ , i.e., food-and-drink apps ( $\sigma_j \approx 0.16$ ) and family apps ( $\sigma_j \approx 0.11$ ), whereas the last used app could not impact the engagement level of next app usage significantly. For example, no matter which app is used before the family apps (as shown in Figure 8d), the time spent on the family apps do not differ much.

**4.3.2 Last Engagement Level.** Besides the last used app, we hypothesize that the time spent on the last app (last engagement level) could also impact how long the user will stay with this app. To focus solely on the impact of the last engagement level, rather than the impact of last used app, we calculate the transition probability between the engagement levels of the last and next usage *within the same app*. We find four typical patterns across all different app categories. We illustrate these correlation patterns between the last and next engagement levels in Figure 9. Across all 45 app categories, 37.8% (17/45) of them follow the pattern of maintaining a higher probability that the next engagement level is as same as the last engagement level (i.e., higher probability in the diagonal of the heatmap), such as the books apps shown in Figure 9a. Besides this common pattern, we also find that 44.4% (20/45) of the app categories have a bidirectional (increasing/decreasing) trend with the closest level, like the food-and-drink apps (i.e., engagement level transition between medium and intensive). The last common pattern is illustrated by comics apps, 13.3% (6/45) of app categories have a similar transition probability across all levels. The remaining pattern is illustrated by travel apps, where no other app categories have similar patterns with them. We can find that there is an apparent increasing trend between engagement level light and medium for travel apps, which could state that users may be more likely to get addicted to planning for a vacation within a travel app.

## 4.4 Long-term Context

**4.4.1 Periodic Pattern.** Some apps are used repeatedly after every specific period. For example, users may check the mail apps every hour or play with a game app every 3 hours [33, 51]. To validate whether the periodic pattern also exists in the app dwell time (i.e., after a specific interval, users may stay with an app again for the similar length of time as before that interval), we first quantify the interval between the two accesses of the same app at the hour level, e.g., 28 min = 0 hour, 1.6 hours = 2 hours. We then generate the histogram of interval time length with different engagement levels respectively for each app category. Two typical patterns are found among all the app categories, which are shown in Figure 10. Figure 10a denotes the interval time length distribution of weather apps, which represents the general trend that a set of most other app categories will obey, i.e., the shorter interval between two accesses of the same app, the less time (engagement level: light) user will spend within the next usage of this app.




Fig. 9. Transition probability of the last engagement level and next engagement level of the same app usage.

However, as shown in Figure 10b (i.e., the shopping apps), the other typical pattern demonstrates that the less time between two accesses, the more likely the second app access will also result in a longer time stay (engagement level: intensive) on the app. This may be because when users are going to buy something, they will browse/revisit the shopping apps multiple times with short breaks (break for checking other information or chat with friends for asking advice) and finally place the order. Additionally, we find that there is a peak around the interval of 24 hours for the engagement level of light/medium for shopping apps. This states that users prefer to spend a similar length of time to regularly check the shopping apps every day (i.e., after every 24 hours), probably for checking the updated discount or product information. This pattern could also be observed with books apps.

**4.4.2 Historical Interest Pattern.** Users' historical interests may be indicative of future intent. For example, users who love sports might potentially access sports apps and consume more sports-related content than others. Therefore it is important to examine the long-term interest patterns of different users. To clarify that if users also have the historical habit for the app dwell time on different app categories, we select three users whose top three preferred (based on usage frequencies) app categories are all the same, which are productivity, widgets, and social apps. We then illustrate the histogram of their historical engagement levels within these three app categories in Figure 11. We can find that *User A* prefers to stay longer (engagement level: intensive) within the widgets apps but spend less time (engagement level: light) with social apps than others. Differently, *User B*



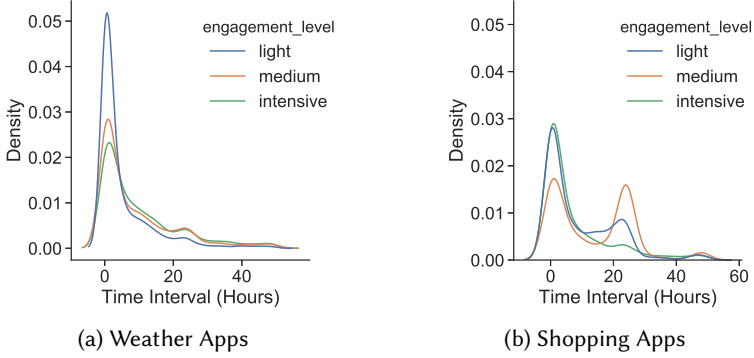


Fig. 10. Two typical interval time distribution with different engagement levels (we limit the x-axis to 50 hours since over 98% intervals are shorter than 50 hours): (1) General trend illustrated by weather apps: the shorter interval between two accesses, the less time is spent with the next usage; (2) Specific daily periodic pattern illustrated by shopping apps: similar length of time is spent on the same shopping app after a specific interval (i.e., 24-hour).

has a longer app usage duration (engagement level: intensive) with social apps. Therefore, even for the users with the same historical app preferences, their engagement habits would be different. In a summary, different users may have their own specific historical engagement patterns on app dwell time for various apps.

### 5 APP USAGE AND ENGAGEMENT PREDICTION

Based on the statistical analysis in Section 4 of app usage duration, we further focus on answering the important question: could we predict which app user will use and how long the user will stay on this app simultaneously? In this section, we start by formulating the next app usage and app dwell time prediction problem, followed by presenting several joint learning strategies for solving these two prediction problems simultaneously.

#### 5.1 Problem Formulation

For the next app prediction, many previous researchers have extracted the predictive features and proposed methods to solve it [26, 49, 63, 70]. Similar to Ricardo et al. [2], we model the prediction of the next app as a supervised classification problem. For the novel prediction problem proposed in our work, which is to infer users’ app dwell time, we also model it as a multi-class classification problem and the continuous dwell time (app usage duration) is represented as the discrete engagement levels illustrated in Section 3.3. More specifically, given the current context  $c$  and a user  $u$ , we need to predict which app  $a$  the user  $u$  will use next and the engagement level  $e$  (light/medium/intensive) for measuring how long the user will stay with this app. We formulate our prediction problem as quaternion, i.e.,  $\{c, u, a, e\}$ . In contrast to a traditional next app usage prediction formulation as  $a^* = \operatorname{argmax}_a \mathcal{F}(a|c, u)$ , we formulate the joint learning task for predicting next app and engagement level as follows:

$$(a, e)^* = \operatorname{argmax}_{a, e} \mathcal{F}((a, e)|c, u). \tag{1}$$

Note that the  $a$  is in the condition for  $e$  since the engagement level is defined based on their corresponding app  $a$  (§ 3.3).




Fig. 11. Historical pattern with app engagement levels.

## 5.2 Joint Learning Prediction Model

We propose three methods to solve the prediction problem formulated in § 5.1: sequential based model (§ 5.2.1), stacking based model (§ 5.2.2) and boosting based model (§ 5.2.3).

**5.2.1 Sequential based Joint Prediction.** The most straightforward method for solving this joint learning problem is to perform these two prediction problems (next app and engagement level prediction) sequentially. It is undoubted that the next app usage prediction is a key issue to the joint learning task. If a predicted app is not the one user will use next, the engagement level prediction might become meaningless. Therefore, we first apply a supervised next app prediction




Fig. 12. Overview of three joint learning strategies: (1) Sequential based joint prediction; (2) Stacking based joint prediction adds a "meta"-classifier to the final stage after we have the prediction results of next app and engagement level sequentially; (3) Boosting based joint prediction has an "error-correction" of next app prediction in the second step for learning of engagement level prediction.

model to measure which app user will use next:

$$a^* = \underset{a}{\operatorname{argmax}} P(a|c, u). \quad (2)$$

Since the engagement level is dependent on which app the user will use next, therefore, another supervised learning model is applied to infer which engagement level the user will have given the predicted app user will use:

$$e^* = \underset{e}{\operatorname{argmax}} P(e|a, c, u). \quad (3)$$

Therefore, the sequential based joint prediction could be formulated as follows:

$$(a, e)^* = \underset{a, e}{\operatorname{argmax}} P(a|c, u)P(e|a, c, u). \quad (4)$$

As shown in Figure 12 (in the green circle), in this sequential-based joint prediction strategy, a next app prediction model is needed to find the proper app  $a^*$  user will use and then an engagement level prediction model is needed to infer the engagement level  $e^*$  based on  $a^*$ . However, the predicted app  $a^*$  is not guaranteed to be the right app; moreover, if the app is not the user will use next, the inferred engagement level based on this app would become meaningless. To this end, we propose two other joint learning methods for predicting these two problems more effectively.

**5.2.2 Stacking based Joint Prediction.** Stacking [50] is an ensemble learning technique that combines multiple classification or regression models via a meta-classifier or a meta-regressor. The base-level models are trained based on a complete training set, then the meta-model is trained on the outputs of the base level models as features. So the main idea for our stacking based joint prediction model is to add a "meta"-classifier to the final stage after we got the prediction results

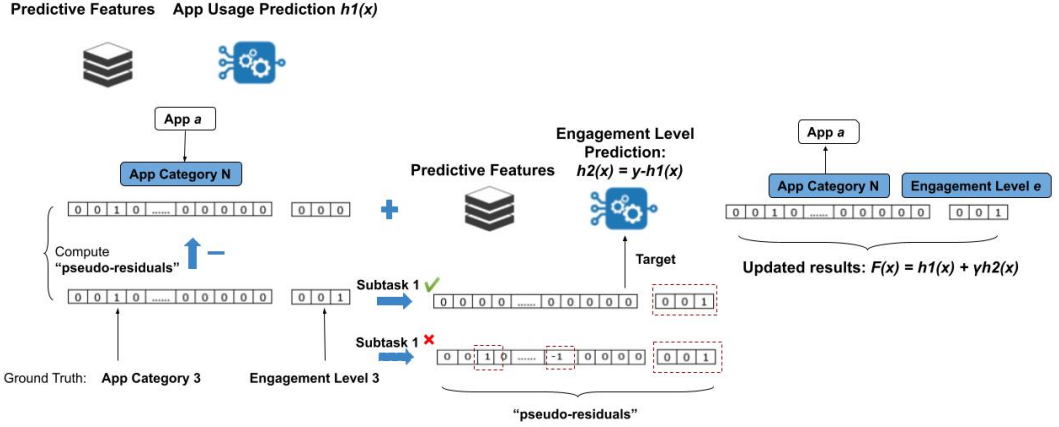


Fig. 13. Boosting based Joint Prediction

on app and engagement level respectively from the sequential based model (§ 5.2.1). Then we may improve the performance by adding this "meta-correction" step at the end of the prediction. In our scenario, the stacking consists of two levels which are base learner as level-0 and stacking model learner as level-1, as shown in Figure 12 (in the orange circle). So the base learners (level-0) are composed of the next app usage prediction model (Eq. (2)) and engagement level prediction model (Eq. (3)), which are the same as in sequential based joint prediction shown in Figure 12. The outputs of each of the sequential classifiers ( $a'$  and  $e'$ ) are collected to create a new dataset. Then the new dataset is used for stacking model learner (level-1) to provide the final output ( $a^*$  and  $e^*$ ). In this way, the predicted classifications from the two base classifiers at level-0 can be used as input variables into a meta-classifier as a stacking model learner, which will attempt to learn from the data on how to combine the predictions from the base models to achieve the best classification accuracy. The stacking based joint prediction in our scenario could be formulated as:

$$(a, e)' = \operatorname{argmax}_{a, e} P(a|c, u)P(e|a, c, u), \quad (5)$$

$$(a, e)^* = \operatorname{argmax}_{a, e} P((a, e)|(a, e)'). \quad (6)$$

**5.2.3 Boosting based Joint Prediction Model.** Boosting [15] is another ensemble method for improving the prediction model of any given learning algorithms. The idea of boosting is to train weak learners sequentially, each trying to correct its predecessor. According to this idea, we aim to fit our two learners (next app and engagement level) iteratively such that the training of the model at a given step depends on the models fitted at the previous steps. Then we can improve the predictions from our first learner of app prediction by adding the "error-correction" prediction in the second learner of engagement level prediction. This is shown in Figure 12 (in the blue circle).

We firstly introduce how boosting method works in the supervised learning problem. For a given data set with  $n$  samples and  $m$  features  $D = \{(x, y)\} (|D| = n, x \in R^m)$ , the goal is to find an approximation  $\hat{F}(x)$  to a function  $\mathcal{F}(x)$  that minimizes the expected value of some specified loss function  $L(y, \mathcal{F}(x))$ . The boosting method assumes a real-valued  $y$  and seeks an approximation in the form of a weighted sum of  $K$  additive functions (called base/weak learners)  $h_k(x)$ :

**Algorithm 1:** Boosting based Joint Training

**Input:** Training data  $D = \{x_i = (c_i, u_i), y_i = (a_i, e_i)\}_{i=1}^n$ .

**Output:** Boosting based joint classifier  $\mathcal{F}(x) = \mathcal{F}((a, e)|c, u) = \sum_{k=1}^K \gamma_k h_k(x)$ .

**Step 1:** Initialize the model with next app prediction

Learn  $h_1(x_i) = P(a_i|c_i, u_i)$  based on  $D$ .

**Step 2:** Improve the predictions from next app prediction by adding "error-correction" prediction in the second learner of engagement level prediction.

1. Compute so-called pseudo-residuals:  $r_i = y_i - h_1(x_i)$ .

2. Fit the second learner  $h_2(x_i)$  to pseudo-residuals. Train it using the training set  $\{(x_i, r_i)\}_{i=1}^n$ .

3. Compute multiplier  $\gamma$  by solving the following one-dimensional optimization problem:

$$\gamma = \operatorname{argmin}_{\sum_{i=1}^n L(y_i, \mathcal{F}_1(x_i) + \gamma h_2(x_i))}.$$

**Step 3:** Update the model:

$$\mathcal{F}(x_i) = \mathcal{F}_1(x_i) + \gamma h_2(x_i).$$

$$\hat{\mathcal{F}}(x) = \sum_{k=1}^K \gamma_k h_k(x). \quad (7)$$

So the model tries to find an approximation  $\mathcal{F}(x)$  that minimizes the loss function  $L$ , and the model is updated as follows.

$$\gamma_k = \operatorname{argmin}_{\gamma} \sum_{i=1}^n L(y_i, \mathcal{F}_{k-1}(x_i) + \gamma h_k(x_i)). \quad (8)$$

To be specific, in our scenario, the dataset  $D$  include features  $x = (c, u)$  and label  $y = (a, e)$ . We have two ( $K = 2$ ) base learners in our prediction, where  $h_1(x)$  is the next app usage learner and  $h_2(x)$  is the engagement level learner. The steps for applying the boosting method in our scenario are as shown in Figure 13 in details:

- (1) Fit a model to the data,  $h_1(x_i) = P(a_i|c_i, u_i)$  for predicting the next app  $a_i$ .
- (2) Compute the pseudo-residuals between the current predicted app  $a'_i$  and the ground truth of finalized results: app with the corresponding engagement level together  $(a_i, e_i)$ . Since the engagement level is defined based on different app categories, additionally, to avoid the sparsity issue caused by using a specific app in comparison, we calculate the "residual" of app prediction results within the app category level. The app category and engagement level are all represented by one-hot vectors and concatenated together during the "residual" calculation as shown in Figure 13.
- (3) Then we fit a model to the residuals,  $h_2(x_i) = y_i - h_1(x_i) = r_i$ . As the boosting based joint prediction process shown in Figure 13, if the first learner output the right predicted app, then the "residual" for the second learner to learn is still only the engagement level ( $r_i = e_i$ ). On the other side, if the first learner output a wrong predicted app, then the residual for the second learner to learn is not only about the engagement level, it also needs to correct the previous results on app prediction.
- (4) Update the new model  $\mathcal{F}(x_i) = \mathcal{F}_1(x_i) + \gamma h_2(x_i)$ . Since the "residual" is calculated based on app category level, the finalized output of app prediction results is also on the category level. We will introduce how to infer the specific app from the predicted app category based on our next app prediction model in § 5.3.

The boosting joint prediction algorithm is shown in Alg. 1. Finally, the boosting based joint prediction could be formulated as:

$$\begin{aligned}
 (a, e)^* &= \gamma_1 h_1(x) + \gamma_2 h_2(x) \\
 &= a' + r \\
 &= a' + [a_r, e] \\
 &= \underset{a}{\operatorname{argmax}} P(a|c, u) + \underset{r}{\operatorname{argmax}} P([a_r, e]|a', c, u),
 \end{aligned} \tag{9}$$

where  $r$  is the "pseudo-residuals" when comparing the current predicted app with the ground truth (next app and engagement level), and  $a_r$  is the difference between the first time predicted app  $a'$  and the target app  $a$ .

### 5.3 Estimating Components within Joint Prediction Strategies

After introducing the three strategies for building the joint prediction model, we can find that the next app prediction and engagement level prediction are two main components for all three strategies. Furthermore, some of these components could be reusable across different strategies, e.g., the next app prediction of sequential-based strategy could also be used by stacking and boosting based strategies at the first stage of app prediction as shown in Figure 12. In the following sections, we discuss how to construct these two components, i.e., the next app prediction (Eq. (2)) and engagement level prediction (Eq. (3)), adaptively for each strategy with our proposed predictive features.

*5.3.1 Next App Usage Prediction.* Many previous researchers [2, 26, 70] have proposed different methodologies to solve the next app usage prediction problem based on personalized mechanism. Additionally, some researchers [14, 67] indicated that the generic (user-independent) model can improve the predictive performance of personalized models when the data is not sufficient. A generic model is trained using data from all available users. Inspired by the work of Do and Gatica-Perez [14], which achieved the best performance by combining the generic model with the personalized model together, we also propose a hybrid next app prediction model with generic and personalized models combined. When building a generic model, the main challenge lies in the fact that the dimensionality of context and output varies depending on all of the users. To learn a generic model and apply it for a given user, generic features and output are needed. In our work, the engagement level is defined based on different app categories. Additionally, the "residual" calculating in boosting based joint prediction strategy (Figure. 13) is also based on app category. Therefore, to ensure the generalization of the generic model and to benefit the further prediction of engagement level, the output of the generic model corresponds to the app category. We infer the specific app user will use based on the user's personalized logs given the predicted app category. Specifically, our proposed hybrid next app prediction model could be formulated as the following function:

$$\begin{aligned}
 a^* &= \underset{a}{\operatorname{argmax}} P(a|c, u) \\
 &= \underset{a}{\operatorname{argmax}} \{P_g(a_c|c, u)P_p(a|c, u), a \in a_c\},
 \end{aligned} \tag{10}$$

where  $P_g(a_c|c, u)$  is the probability that user  $u$  will use app category  $a_c$  based on our generic app category prediction model; and  $P_p(a|c, u)$  is the probability that user  $u$  will use the app  $a$  based on their own personalized app prediction model and we will limit the  $a$  to the apps belonged to the predicted app category  $a_c$ . It means we firstly get the app category  $a_c$  user will use based on all users' logs and then further rank the apps belonging to this app category based on the personalized

logs of user  $u$ . Through this way, we can have the intermediate output of the app category, which could be used for benefiting the later engagement level prediction. What is more, this generic category-level prediction model could alleviate the cold-start problems with new apps/users resulted from user-specific prediction models, which may be trained on a limited quantity of logs.

For the generic app category prediction model, we extracted all the features that have correlation with app usage patterns, which have been validated by previous works [14, 67] and also available in our dataset (§ 3.1): User characteristics (age, gender, country, device type, operation system) [29, 31, 68], temporal features (hour of day, day of week) [26, 33, 49], historical preferences [14], last one/two apps used [70] and periodic features [33, 35, 51]. To enrich our predictive features, we further expand users' characteristics by adding their total app usage duration, total app usage frequency, and unique app amount as user characteristic features. To ensure the generalization of the generic model, users' historical preferences are represented based on the total access frequency of each app category. Moreover, we also add users' app preferences on the last day, the last hour and the last session to expand the context features. In summary, we extracted 14 features related to app usage patterns from 4 aspects: user characteristics, temporal features, short-term context features, and long-term context features. All the features used for next app prediction are described in Table 7. Then the personalized next app prediction model is generated based on the predictive features proposed in previous works [2, 49, 70]: hour of day, day of weekday, most recently used apps (the last one/two apps), and periodic feature.

Table 7. All the features used in our next app prediction model related to users characteristics and context

Feature Type	Feature	Description
<b>User and Device</b>	Age*	Users' age group: 13-17, 18-24, 25-34, 35-54 and 55+
	Gender*	Users' gender: male and female
	Device Type*	Users' device type: phone and tablet
	Total App Usage Duration*	Users' total app usage duration for all apps
	Total App Usage Frequency*	Users' total access frequency for all apps
	Total Unique App Amount*	The amount of unique apps the user has accessed
<b>Context</b>	Temporal Features	
	Hour of the Day*	Different hours of a day: 0 - 23
	Day of the Week*	Different days of a week: Monday to Sunday
	Short-term Context Features	
	App Preference in the Last Day	Access frequency of each app category in the last day
	App Preference in the Last Hour	Access frequency of each app category in the last hour
	App Preference in the Last Session	Access frequency of each app category in the last session
	Last Used App*	The last used app in the same session
	Last Used Two Apps	The last used two apps in the same session
	Long-term context Features	
	Periodic Feature	Time intervals between the current time and the last usage of each app category
Historical App Preference	Total access frequency for each app category	

\*: The features marked with \* are the common predictive features also used in app engagement level prediction.

**5.3.2 App Engagement Level Prediction.** In this subsection, we further explore the app engagement level prediction models that could fit in the different joint prediction strategies. As we mentioned before, we model the novel prediction problem of app engagement level as a multi-class classification problem, where the engagement level is defined based on different app categories. In the sequential based joint prediction strategy (§ 5.2.1) and the level-0 classifiers of the stacking based joint prediction strategy (§ 5.2.2) (Figure 12), the engagement level prediction model could be trained for each app category respectively. For example, if we have predicted that the user will use the news app  $a^*$  next, we just select the classifier which has been trained specifically based on the logs of using news apps. Then we use this classifier to predict the engagement level. So the engagement level prediction model for sequential-based and stacking-based joint learning strategies could be

formulated as:

$$\begin{aligned}
 e^* &= \underset{e}{\operatorname{argmax}} P(e|a, c, u) \\
 &= \underset{e}{\operatorname{argmax}} \{P_{a_c}(e|c, u), a^* \in a_c\},
 \end{aligned} \tag{11}$$

where  $a_c$  is the corresponding app category of the predicted app  $a^*$  based on the next app prediction model. It states that we select the engagement level classifier exactly based on the prediction results coming from the next app prediction results. For each app category, we extract the predictive features based on the analysis in section 4: demographic features, device features, hour of day, day of week, last used app, last engagement level, last engagement level of the same app category, periodic feature and historical engagement level feature. Besides the common user characteristic features and some of the context features that have been listed in Table 7, we describe the additional features for the app engagement level prediction model in Table 8. These features are all established to have impacts on users' app dwell time in § 4, and we will further analyze the feature importance in the following experimental results section § 6.4.1.

Different from the engagement level prediction model  $P(e|a, c, u)$  in sequential (Eq. (4)) and stacking (Eq. (5)) based strategies, the boosting strategy need to infer the "residual" with  $P(r|a, c, u)$  (Eq. (11)) and then sum it to the first app prediction result for getting the finalized app and engagement level together. During this process, the next app user will use would be re-inferred during the engagement level prediction, we cannot select a specific classifier based on the first predicted app. Therefore the engagement level prediction model of the boosting-based strategy can only be constructed as a generic model which could be applied to all apps instead of a specific app category. Specifically, the first predicted app  $a'$  from the next app prediction model would be treated as an input feature in the following engagement level prediction, which is represented as a one-hot vector as shown in Figure 13. Additionally, to ensure the generalization of the input in the engagement level prediction model within boosting-based joint prediction, all the predictive features will also need to be expanded for representing the behaviour pattern coming from all the different app categories (e.g., the feature of historical engagement level for predicted app category need to be extended to historical engagement level for all different app categories). In this model, no matter which app category is predicted to be used next, the same engagement level prediction model will be applied. Since the finalized app and engagement level would be inferred together within the boosting strategy, the formulation of engagement level prediction could not be decomposed, which has been shown in § 5.2.3.

Table 8. Additional features used in our app engagement level prediction models related to user characteristics and context

Feature Type	Feature	Description
Context	Short-term Context Features	
	Last Engagement Level	The last engagement level of all app categories
	Last Engagement Level of Predicted App Category	The last engagement level of the usage on predicted app category
	Long-term Context Features	
	Periodic Feature	Time intervals since the last use of all app categories
	Periodic Feature of Predicted App Category	Time intervals since the last use of predicted app category
	Historical Engagement Levels	Historical sum of engagement levels for all app categories
	Historical Engagement Levels of Predicted App Category	Historical count of each engagement level for predicted app categories



## 6 EXPERIMENTAL RESULTS

In this section, we measure the performance of our proposed prediction models (§5). Experiments are conducted on a real-world app usage log data (§3.1). We use 70% of the data for each user as training data and the remaining 30% as test data. We first evaluate the performance of our proposed prediction model on the classic prediction problem, predicting the next app (§5.3.1). Then the three joint learning strategies for predicting the next app and engagement level simultaneously (§5.2) are thoroughly evaluated from different perspectives.

### 6.1 Evaluation

In our proposed prediction models, we measure the performance of all prediction problems based on four metrics: accuracy, precision, recall and F1 score: The accuracy in our problem is defined as the fraction of correctly classified samples; The precision is defined as the ratio:

$$\frac{tp}{(tp + fp)}, \quad (12)$$

where  $tp$  is the number of true positives and  $fp$  the number of false positives. The precision is intuitively the ability of the classifier not to label as positive a sample that is negative. The recall is the ratio:

$$\frac{tp}{tp + fn}, \quad (13)$$

where  $fn$  is the number of false negatives. The recall is intuitively the ability of the classifier to find all the positive samples; The F1 score can be interpreted as a weighted average of the precision and recall, where an F1 score reaches its best value at 1 and worst score at 0. The relative contribution of precision and recall to the F1 score is equal. The formula for the F1 score is:

$$F1 = 2 * \frac{\text{precision} * \text{accuracy}}{\text{precision} + \text{accuracy}}. \quad (14)$$

Please note that in our multi-class case, all the precision, recall, and f1 scores are the weighted averages of scores for each class. In order to account for label imbalance, metrics are calculated for each label, and their averages are weighted by support (the number of true instances for each label).

For the next app prediction results, we will count the item as correctly predicted when the predicted app is exactly the same as the ground truth app. While evaluating the joint prediction problem (both next app and engagement level), a correctly classified sample indicates that both the predicted app and predicted engagement level are correct. If either of the prediction is wrong, it will not be counted as a correct classification.

### 6.2 Baselines

*6.2.1 Next App Usage Prediction.* In order to comprehensively measure the performance of our proposed hybrid next app prediction model (§5.3.1), we compare it with state-of-the-art counterparts. Based on the available sources of evidence in our dataset, we first select the two common baseline methodologies: MFU (Most Frequently Used) and MRU (Most Recently Used) [49, 62, 70]. We then tested additional three methodologies from previous works which conducted the next app prediction by combining app usage history and contexts in a unified manner: SVM+Context [49], CPD [51] and BN [70]. Recently, the neural approaches are popular in solving the time sequence problems, we also investigated the performance of app usage prediction based on LSTM [61] as the baseline:

- *MFU*: the predicted app is the most frequently used app.

- *MRU*: the predicted app is the most recently used app.
- *SVM+Context*: Shin et al. [49] used a SVM classifier [4] and context information (i.e., day of week, hour of day, last used app and time since last app usage) to predict the next app user will use.
- *CPD*: Tan, et al. [51] proposed a prediction framework: Prediction Algorithm with Fixed Cycle Length (PAFCL). They hypothesized that each application has different usage probabilities in the different time slots of a cycle. CPD (Cumulative Probability Distribution) is the method used to choose applications with higher probability as the candidates based on computing the probabilities of each used app in the specific time slot.
- *BN*: Zou et al. [70] proposed a Bayes Network (BN) model which is a linear combination of  $p(a_n = A|a_{n-1} = A_{n-1})$  (based on last used app  $A_{n-1}$ ) and  $p(a_n = A|a_{n-2} = A_{n-2})$  (second last used app  $A_{n-2}$ ).
- *LSTM*: Xu et al. [61] proposed a generic prediction model based on Long Short-term Memory (LSTM), which is an enhancement of the recurrent neural network (RNN) model. The proposed model converts the temporal-sequence dependency and contextual information into a unified feature representation for the next app prediction.

*6.2.2 Next App and Engagement Level Joint Prediction.* Since we are the first work proposed to predict which app user will use and how long the user will stay on that app simultaneously, where the time spent has a high dependency with which app user is engaging, no previous related work could be identified as our baseline methodologies. Therefore, we propose the baselines from two kinds of approaches. Firstly, by following the classic baselines (MFU and MRU) in the next app prediction problem, we present two naive prediction methods as the baselines for our proposed joint prediction problem. Secondly, two recent works on dwell time prediction of online services (e.g., video [59] and media streaming [57]) are selected as baselines. Similar to our work, how long a user stays with an item in these services originally have associations with the content category, and could also be affected by user characteristics and contextual features. The only difference is that they solely predict how long a user will stay based on the specific item, whereas no prediction on which item the user will engage is required. In order to make them comparable baselines, we assumed the ground truth of the next app is known and leveraged those two prior works [57, 59] for predicting engagement (dwell time) of the oracle app. This demonstrates the upper bound of those approaches (oracle performance) within the joint prediction task.

- *MFU*: In our joint prediction problem, we discuss the usage frequency of the two fields together: the next app and engagement level. Then the MFU baseline states that every time we recommend the tuple  $(a, e)$  as the app  $a$  and engagement level  $e$  based on the popularity of the tuples.
- *MRU*: Similar to the MRU in next app prediction, we hypothesize that a correlation between two sequential tuples of app and engagement level may exist. For example, a user may always like to access the weather app for 10 seconds and then access the news app for 5 minutes. Therefore we generate a correlation-based baseline approach for our joint prediction problem, which aims to predict the next app and engagement level based on the correlation between two sequential tuples  $(a', e')$  and  $(a, e)$ . We first calculate all the transition probabilities from one tuple to another. When we know the last usage tuple is  $(a', e')$ , we could recommend the tuple  $(a, e)$  that has the highest probability to be used next.
- *CSP*: Wu et al. [59] conducted a large-scale measurement study of engagement on videos. They predicted engagement from video context, topics, and channel reputation, etc. They used linear regression with L2-regularization to predict engagement metrics and state that the Channel Specific Predictor (CSP) performs best, which is to train a separate predictor for

Table 9. Performance comparison of next app prediction models (\* indicates statistical significant ( $p \leq 0.01$ ) using two-tailed T-test when compared our hybrid next app prediction model to the best baseline model (LSTM).

Method	Measurement			
	Accuracy	Precision	Recall	F1
MFU	0.486	0.489	0.486	0.486
MRU	0.518	0.521	0.518	0.518
CPD [51]	0.424	0.419	0.424	0.369
BN [70]	0.453	0.483	0.453	0.467
SVM+Context [49]	0.606	0.563	0.606	0.572
LSTM [61]	0.525	0.660	0.525	0.576
Our Hybrid Model *	<b>0.640</b>	<b>0.607</b>	<b>0.640</b>	<b>0.613</b>

each channel instead of using the shared predictor. To fit into our scenario, we generate the same CSP for each app category with their proposed features available in our dataset (e.g., the one-hot encoding of category, mean number of daily usage, mean, std and five points summary of past engagement (dwell time), etc.). The performance of the joint prediction problem is reported based on the oracle app and the predicted dwell time (which would be transformed to engagement levels according to our definition in Sec. 3.3.)

- *Aggregated*: Vasiloudis et al. [57] presented the first analysis of session length in a mobile-focused online service (i.e. music streaming service). They showed that the time length of sessions can differ significantly between users. They used gradient boosted trees with appropriate objectives to predict the length of a session using contextual and user-based features. Their experiment results showed that the aggregated model trained with all the data performed better than the personalized models trained on each user’s data. To fit into our prediction problem, we also trained the aggregated model based on all data with their proposed features available in our dataset (e.g., gender, age, device, duration of the user’s last session and time elapsed since the last session, etc.). As we mentioned before, the performance of the joint prediction problem is reported based on the oracle app and the predicted dwell time (which would be transformed to engagement levels according to our definition in Sec. 3.3).

### 6.3 Hybrid Next App Prediction Model

Our analysis starts with the hybrid next app prediction results given it is the unified component leveraged by all the three different joint learning strategies. It is used at the first step for inferring the next app user will use before predicting its engagement level (Figure 12). Table 9 shows the performance of baseline methodologies and our proposed hybrid next app prediction model. We test our model with a set of state-of-the-art classification models, including Random Forests [7], L2-regularized Logistic Regression [24], K Nearest Neighbours [11] and Support Vector Machines [8]. These models construct different prediction functions for the data from different aspects, and they can provide more robust results for our prediction. Since our proposed hybrid next app prediction model (§5.3.1) involves the prediction results of two predictive components (Eq. (13)), we select the best classifiers for both of them within the hybrid model, which combines the results of generic app category prediction with Random Forest classifier and the results of personalized next app prediction with SVM classifier.

From Table 9, we can find that our proposed hybrid next app prediction model could significantly improve the performance of the best baseline model (LSTM) by 6.4% on the F1 measure. The worse performance of LSTM is expected since it could not incorporate any user characteristics

Table 10. Performance comparison of joint learning prediction models (\* indicates statistical significant ( $p \leq 0.01$ ) using two-tailed T-test compared to the best baseline (CSP)).

Model	Measurement			
	Accuracy	Precision	Recall	F1
MFU	0.286	0.286	0.286	0.286
MRU	0.308	0.307	0.308	0.308
Aggregated [57] *	0.347	0.742	0.347	0.448
CSP [59] <sup>†</sup>	0.339	0.729	0.339	0.467
Sequential*	0.375	0.382	0.375	0.369
Stacking*	0.374	0.381	0.374	0.365
Boosting*	<b>0.485</b>	<b>0.483</b>	<b>0.485</b>	<b>0.483</b>

<sup>†</sup>: We assumed the ground truth of the next app is known and leveraged these baselines for predicting engagement (dwell time) of the oracle app. The performance reported in table demonstrates the upper bound of those approaches (oracle performance) regarding the joint prediction problem.

and contexts (e.g., access time, device type, etc.) during the prediction. While our proposed hybrid next app prediction model not only takes the user characteristics and contextual information into consideration but also learn from users' common patterns to overcome important sources of prediction errors resulting from insufficient training data. It also states that through the use of community similarity and common usage patterns learned based on different app categories, we can improve the next app usage prediction by identifying the generic usage patterns present in similar users - rather than relying solely on the specific app usage patterns. This is also consistent with the findings from previous studies [14, 67] which claim that the generic model can improve the predictive performance of models solely based on individual's logs. Therefore, our proposed hybrid next app prediction model is adopted to apply to further engagement level prediction for different joint learning strategies.

#### 6.4 Next App Usage and App Engagement Level Joint Prediction Strategies

We first evaluate how different models perform for solving our proposed joint prediction problem: which app user will use next and how long the user will stay with this app? Then we further investigate the predictive ability of features in the two prediction problems respectively and the prediction effectiveness of our proposed different joint prediction strategies.

As mentioned, our hybrid next app prediction model is applied to the first stage prediction on the next app for all the joint learning strategies (Figure 12). For the prediction of engagement level, the predicted app will be represented in different ways based on different joint prediction strategies. For the sequential and stacking based joint model, the predicted app will be used to select the specific engagement level classifier with the corresponding app category (Eq. (14)). For the boosting based joint model, the predicted app coming from the first stage of next app prediction could only be treated as the input feature for the next step prediction of engagement level (Eq. (12)). Table 10 shows the ultimate performance of different joint prediction strategies and all baselines.

We can find that all our proposed joint prediction models are better than the two classic baselines: MFU and MRU. For another two baselines, we can observe that even we assumed the ground truth of the next app is known, the upper bound performance of these approaches could not beat our proposed best joint prediction model based on boosting strategy. To be specific, it states that we assume the predicted app is exactly the same as the ground truth and only evaluate the engagement level prediction performance, these baselines have performed worse than our boosting-based joint prediction model. It mostly because they didn't model the engagement with comprehensive characteristics as our proposed model, where we take the user characteristics,

Table 11. Feature importance (MDI) of app category prediction.

Feature	Feature Type	Importance (MDI)
Total_App_Usage_Frequency	User	0.023
Hour	Temporal	0.011
Periodic_Feature	Long-term	0.011
Total_App_Usage_Duration	User	0.010
Historical_App_Preference	Long-term	0.005
Total_Unique_App_Amount	User	0.005
Age	User	0.004
App_Preference_Last_Day	Short-term	0.003
Weekday	Temporal	0.002
Gender	User	0.002
App_Preference_Last_Hour	Short-term	0.002
Device_Type	User	0.001
App_Preference_Last_Session	Short-term	0.001

short/long-term usage patterns all into consideration. Therefore, if the next app prediction task is added, the worse performance of the joint prediction problem should also be expected for these baselines. Among all the three proposed joint prediction strategies, stacking and boosting based strategies are two advanced approaches originally proposed to improve the performance with the most straightforward strategy, sequential based strategy. However, we find that the stacking based strategy does not improve the performance when compared with the benchmark sequential based strategy. This might be due to that, in our scenario, the base learners within stacking are not trained for the same target (we have two different prediction tasks: app and engagement level) where the stacking principles do not apply. On the other hand, the boosting strategy works best compared to all other joint models. It respectively outperforms the baseline models over 56% and the sequential/stacking models about 31% on F1 measure. We mainly focus on investigating the boosting and sequential strategies in the later sections about how boosting helps in the joint prediction. Firstly, besides the overall performance reported in Table 10, we look into the detailed prediction results. It demonstrates that the accuracy of next app prediction and engagement level prediction respectively are: 64% (app) and 58.6% (engagement level) for sequential strategy, 85.6% (app) and 56.7% (engagement level) for boosting strategy, where the accuracy of engagement level is calculated only based on the data with right predicted apps. So we can find that the improvement of overall performance for boosting strategy is mainly because of the "error-correction" step which corrects the app prediction results.

In the following sections, we will first analyze the feature importance within the two prediction problems respectively, especially focus on exploring the predictive ability of the newly proposed features. Then we will dig into how is the effectiveness of boosting based strategy compared to the sequential based strategy.

**6.4.1 Feature Analysis.** The sequential based strategy is implemented by conducting the app prediction and the engagement level prediction sequentially, where the app prediction is absolutely independent with the further engagement level prediction. Therefore, we discuss the analysis of the most impactful features for those two tasks respectively. The next app prediction within the sequential strategy is implemented as same as our proposed hybrid next app prediction model, which would infer what app category the user will use and then select the specific app given this predicted app category. To make the comparison more intuitively, we conduct the analysis within the app category level. As we mentioned above, the random forest is selected as the best classifier for the app category prediction problem. The random forest can be used to rank features by their importance in the classifier, which provides useful insights about the discriminative power of the features in the considered problem setting. *Mean Impurity Decrease (MDI)* is the most common

Table 12. Top feature weights (standardized coefficients  $\geq 0.001$ ) for logistic regression model of engagement level prediction. \* indicates p-value  $\leq 0.001$  using Chi-Squared test.

Feature	Feature Type	Weight
Historical_Level_Light*	Long-term	0.571
Historical_Level_Medium*	Long-term	0.197
Total_App_Usage_Duration*	User	0.175
Historical_Level_Intensive*	Long-term	0.112
Periodic_Feature*	Long-term	0.070
Age*	User	0.032
Last_Used_App*	Short-term	0.019
Hour*	Temporal	0.008
Weekday	Temporal	0.002
Last_Engagement_Level*	Short-term	0.001

way to obtain feature importance from random trees [7]. It is computed by averaging across all the trees in the forest the amount of impurity removed by each feature while traversing down the tree, weighted by the proportion of samples that reached that node during training. Using this method, we obtained the feature importance for all features in the app category prediction, as listed in Table 11.

We find that besides the critical temporal feature *hour*, the most important categories of features are mostly user and long-term context features. Compared to user characteristics features such as demographics and devices, the total usage (frequency, duration, and unique app amount) maintain higher impacts on the next app category prediction. The hour of day feature has been used to identify the salient pattern within app usage behaviour in many previous works [26, 33, 49], e.g. the user usually set an alarm at around 23:00. The popularity of app usage (historical app preference) is also a famous feature that has been established by previous works [14, 35]. Additionally, our finding of the periodic pattern is consistent with the previous works on next app prediction, where the periodic patterns have been identified as the effective feature for inferring the app usage pattern [33, 35, 51]. Liao et al. [35] also stated that the periodical usage feature is the most difficult one to be substituted by other features.

The engagement level prediction is the novel problem proposed in our work and we extracted many different predictive features from user&device characteristics, temporal pattern, and short/long-term context (§4) to infer how long the user will stay with an app. Similarly, since the engagement level prediction within the sequential strategy can be studied separately, we discuss the impacts brought by different features for engagement level prediction within the sequential based strategy. We opt to build a classifier for predicting users' engagement with three levels: *light*, *Medium* and *Intensive*. Hence, the problem of modelling user engagement turns into a multi-class classification problem. Similar to the next app prediction, we test and empirically compare the performance of a wide range of classification techniques, including Random Forests (RF), L2-regularized Logistic Regression (LR), K Nearest Neighbours (KNN), and Support Vector Machines (SVM), for predicting the app engagement level, where the Logistic Regression classifier performs best.

Then we examine the contribution of each feature based on the feature coefficients in the logistic regression model. To compare the importance of different features, we divide each numeric variable by two times its standard deviation [20]. Through this way, the resulting coefficients are directly comparable for both binary variables (e.g., categorical dummy variable) and numerical features. Table 12 reports the top feature weights (standardized coefficients  $\geq 0.001$ ) of the Logistic Regression model for the engagement level prediction. We can find that most of the top features are originated from the long-term context features: historical engagement level preference and periodic pattern. It is not surprising that the historical pattern has more influence on how long a user will stay with an app. If the user always prefers to play games for a longer time, then he may still spend more




Fig. 14. The confusion matrix of our prediction model on the engagement level.

time in the game app this time. The periodic feature has more impacts on users’ engagement level than the temporal context, hour, and weekday. This demonstrates that no matter when the user uses this app, the time since the last use of this app is more important for inferring how long the user will engage with this app. Additionally, the last used app has more impacts on predicting the engagement level compared to all short-term context features. This could provide more insights for the app developers to recommend the contents of different time lengths according to the last used app to improve the user experience. Another finding is that we observe age is the most important signal among all demographics and device characteristics when inferring the app usage duration.

**6.4.2 How Effective are Boosting-based vs. Sequential-based Strategies?** We have shown the boosting based strategy outperforms the benchmark sequential based strategy by the performance margin of 31% on the F1 measure. Therefore we further conduct some error analysis to understand the underlying reasons.

Firstly, for the novel prediction problem on engagement level, the sequential and boosting based strategies achieve similar performance, which are 58.6% and 56.7% respectively on accuracy. It demonstrates that even the boosting based strategy gets better overall performance in our joint prediction problem, it cannot improve the engagement level prediction results. We further analyze the prediction ability on app engagement level prediction with our proposed sequential strategy. We select only the engagement level prediction results when the next app is correctly predicted. Figure 14 shows the confusion matrix of the app engagement level prediction results. For the wrongly predicted results of all engagement levels, we can find that they have higher probabilities to be predicted into the adjacent level. For example, for the intensive engagement level, 19% of them are misclassified as a medium level, which covers about 95% of the wrongly predicted results. Similarly, for the light engagement level, 25% of them are misclassified as a medium level, which is higher than those that are misclassified as intensive. The engagement level prediction results of boosting based strategy also get a similar confusion matrix as shown in Figure 14. We now focus on exploring how boosting is more effective in the next app prediction as follows.

For the next app prediction, we have reported that the boosting based strategy improves the accuracy from 64% to 85.6%, which is resulted from the "error-correction" step. We then compare the app prediction results of sequential based and boosting based strategies. For the sequential based strategy, we found that the incorrect results are mainly generated from three reasons:

(a) **Global popularity.** We find that for 66.7% (30/45) of the app categories, the top category they were misclassified as is *productivity*. For example, 37.4% of widgets apps are wrongly predicted

Table 13. Illustration of Case Studies in Next App Prediction

Error Reason Type	Predicted Result	Ground Truth	Correction while Boosting	Correction/Uncorrection Reason
Global Popularity	productivity	widgets, medical, transportation	widgets, medical, transportation	The "pseudo-residuals" are applied as ground truth in boosting makes the model less impacted by the global popularity of app categories within the dataset.
Short Usage Interval	card	casino	casino	Additional information are provided by the engagement features added in the boosting strategy. Last engagement of casino is longer than the card app, so these cases are corrected from card to casino.
	arcade	adventure	No	No additional information could be provided based on engagement.
Similar App Usage Frequency	lifestyle	food-and-drinks	food-and-drinks	Boosting helps when the engagement pattern provides additional information. Specifically, the results are corrected by boosting since the historical engagement pattern of food-and-drink apps usage of this user is longer than lifestyle apps.
	health-and-fitness	medical	No	None of them are corrected by boosting since no difference exists in the historical engagement pattern of these two apps of the user.

as productivity, 31.7% of medical apps are wrongly predicted as productivity, and 27.2% of transportation apps are also wrongly predicted as productivity. This could be easily explained since the productivity apps have the highest global popularity among all the app categories. It results in our generic app category classifier inferring that the user has a higher probability to use the productivity apps given all users' data. By analysing these cases, we find that the boosting strategy resolves most of the issues resulted from global popularity. For example, 25.4% of widgets apps and 12.8% of transportation apps are corrected by boosting from being misclassified into the productivity apps. This is because that the prediction model in the boosting step does not involve the global popularity of app categories. Since the "pseudo-residuals" (Figure 13) are applied as the ground truth, this makes the model more sensitive for predicting users' app usage based on the context and user characteristics, rather than the global popularity of app categories within the dataset.

(b) **Short usage interval.** Another important factor that accounts for the misclassification within the sequential based strategy is the time interval (periodic feature) between app usage. We observe that for those users whose historical app preference deviates from an average user (e.g., productivity is not the top preferred app), the periodic feature (time since the last usage) would have stronger impacts. For example, 30.8% of the adventure apps are wrongly predicted as arcade apps given the time from the last usage of arcade app is short. Since the model learned that the shorter the time interval, the higher probability the same app would be used again. In these cases, even the adventure has a higher usage frequency in the user's historical app usage pattern, it would be wrongly predicted as arcade since arcade has a higher probability to be re-accessed in a short period of time.

(c) **Similar app usage frequency.** The third major reason for the misclassified cases is that when two app categories are all recently used with a similar frequency, the one with the higher personalized historical usage frequency would be selected. For example, 11.4% of medical apps are wrongly predicted as health-and-fitness apps because that the medical app and health-and-fitness app are always used in the same session, the health-and-fitness app is predicted as the result since it has higher popularity within the user's historical app usage pattern.

For the latter two error analysis in the sequential based strategy (described above as (b) and (c)), we find that the boosting strategy only helps some of the cases when the engagement information could provide additional insights. For example, among the wrongly predicted 30.8% adventure apps, only 5% of adventure apps are corrected from arcade by boosting since the last engagement



level of adventure is longer than arcade. For those cases without this additional information from engagement, they are still wrongly predicted. However, for some specific cases, e.g., 7.8% of casino apps are predicted as card apps for the same reason as adventure&arcade apps, but all of these 7.8% cases are corrected to casino with boosting. This is because for all the cases the last engagement of casino is longer than the card app. Similarly, for the cases which are wrongly predicted because they have similar recent usage patterns, the boosting only helps when the engagement pattern provides additional info. For example, 3.7% of food-and-drink apps are corrected by boosting from lifestyle apps. This is because the historical engagement pattern of food-and-drink apps usage of this user is longer than lifestyle apps. However, for the 11.4% wrongly predicted medical apps, none of them is corrected by boosting since no difference exists in the historical engagement level pattern of these two apps.

To make the benefits brought by boosting strategy clearer, we also added a summary table of the case studies as Table 13, where we can find the corresponding reasons about why some of the cases are corrected by boosting but some others are not.

## 7 DISCUSSION

In this study, we explored the factors that affect users' app dwell time from users' app usage logs and show evidence that the next app and how long the user will stay on this app could be predicted simultaneously. First, to answer our first research question (**RQ1**), we take a systematic approach to uncover the dependency of users' app usage duration on user characteristics and context features based on a large-scale dataset. We then showed that the features related to users' historical engagement pattern and periodic usage pattern are good predictors of how long users will stay with an app. To solve the second research question (**RQ2**), we propose three different joint prediction strategies and demonstrate that the boosting based strategy performs the best. We further conduct the error analysis on the boosting strategy compared to the benchmark sequential based strategy. Based on the analysis, we find that besides the benefits brought by the "pseudo-residuals" within the boosting principle, the engagement features also provide additional insights to help infer which app user will use. These findings inspire us that we should think of adding more engagement related features when predicting the next app. As below, We discuss how our findings can be applied to future mobile systems and applications, the limitations in our study and the differences between personalized and general engagement prediction.

### 7.1 Implications

Firstly, the model proposed in our work can be applied for more tailored engagement-aware recommendations on mobile phones. As the operating system has access to all the features used for training the next app and app engagement models in this work, it is uniquely suited to predict a user's likely next app usage and engagement level under the current context. By doing so, the operating system can manage the delivery of content and services to the end-user by matching their engagement demands with the predicted engagement levels of the user. For instance, an app that shows mobile advertisements require high engagement from the end-users and can ask the app provider to push its content to the user when he/she is likely to be highly engaged. The media apps, like video and news apps, could recommend more satisfactory content based on the predicted engagement level of users to improve user experience.

Recently, there is an active area of research for the timely delivery of notifications on mobile devices. Researchers have primarily focused on understanding the receptivity of mobile notifications [37] and predicting opportune moments to deliver notifications in order to optimize metrics such as response time [41]. While response time is indeed a useful metric to optimize for, they do not capture how much engagement the user will show towards the notification. The primary purpose

of a notification is to attract user attention and increase the possibility of user engagement with the notification content. As such, we believe that the models and features we explored in this work can also be incorporated in designing an effective notification delivery mechanism.

## 7.2 Limitations

Although we have explored several meaningful features that could be applied in benefiting users' app engagement (dwell time) prediction, we must acknowledge the limitations of the dataset used in our study. Firstly, similar to all other real-world app usage datasets, the app popularity distribution follows Zipf's law [32, 42], which indicates that only a few apps have high installation/usage whereas many apps have low installation/usage. But doing a balanced prediction would not fit the realistic evaluation settings, which would bring the bias of the users/contexts to be selected. To handle this imbalance issue within app categories while evaluation, we measure the performance of all prediction problems based on four metrics: accuracy, precision, recall and f1. All the metrics are calculated for each label, and their averages are weighted by support (the number of true instances for each label). We also conducted case studies to explore the prediction results based on different app categories. Additionally, our dataset might not be representative of the entire population of mobile users, as the users and apps are only coming from the apps registered in this library. This means that not all the apps usage behaviour of users could be tracked and there may be a selection bias in the subset of users being studied. While this might occur to some extent, given the scale of our dataset (with over 1.3 million logs), we believe our data would still provide useful insights, and our predictive models are the best models so far, effective for most users.

Lastly, the predictive features we extracted can be further enriched. Other features that precisely characterize users' app engagement behaviours could be further explored, such as the location, wifi access status, battery, device mode setting (e.g., silent), illumination, screen, and blue-tooth. In this work, our main aim is to validate that how long a user will stay with an app could be modelled based on the user characteristics and context features. We would like to further explore additional features and models for improvements in our future work.

Despite these limitations, we believe our work is the first-of-its-kind study to examine, and model the mobile app engagement purely based on features derived from users' app usage logs. We also hope that the framework, models, and insights developed in this work can bring clarity and guidance to aid future mobile system developers in designing better, and engagement-aware user experience.

## 7.3 Personalized V.S. General Engagement Prediction

We know that the performance of the personalized model could be highly impacted by whether there is sufficient data for training or not. In this work, we are interested in relatively short-term user app engagement patterns. Therefore, we collect app usage data from all users for a week. Due to the nature of the data, it might be more difficult to acquire sufficient per-user patterns, which is also validated by our experiment results regarding the next app prediction (i.e. personalized model could not outperform our proposed hybrid model). When the long-term data is available, the personalized approach might be more suitable and achieve better performance. We leave the exploration of the personalized model for app usage prediction in our future work. In terms of the app engagement level (dwell time) prediction, we introduce the engagement levels according to different app categories for handling users' different consumption behaviour on different contents. Hence our engagement level prediction results do not directly translate to user-specific engagement. Additionally, we believe that defining and measuring aggregate engagement is also useful for content producers, e.g., video producers on Youtube. The content providers often do not target a specific user, but a large number of audience. Other than predicting how long user will stay on

the app, we also want to provide content producers with a new set of tools to create engaging content and forecast user behaviour. For future work, we would measure users' engagement at the personalized level as complementary to the aggregated engagement study, especially when we have more sufficient training data for individuals. It would help the mobile apps to provide more fine-grained services to a specific user based on the more accurate expected time length.

## 8 CONCLUSION

In this paper, we propose to predict both users' next app and the engagement level at the same time. For the first time - to the best of our knowledge - a comprehensive analysis of users' app dwell time is conducted based on the large-scale commercial mobile logs, especially focusing on inferring the correlations between different predictive features and users' app usage duration. We find that the users' historical engagement pattern, periodic behaviour pattern, and the recent usage pattern have more impacts when inferring users' app dwell time. To solve our joint prediction problem on the next app and app engagement level, we propose three strategies, where the boosting based joint prediction model works best. Our experimental results demonstrate that users' next app and engagement level could be effectively predicted at the same time, and our proposed prediction method outperforms all baseline experiments by a large margin. Our work can help for providing more satisfying services to users for improving users' experience on mobile devices.

## REFERENCES

- [1] Eugene Agichtein, Eric Brill, and Susan Dumais. 2006. Improving web search ranking by incorporating user behavior information. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 19–26.
- [2] Ricardo Baeza-Yates, Di Jiang, Fabrizio Silvestri, and Beverly Harrison. 2015. Predicting the next app that you are going to use. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM, 285–294.
- [3] Mikhail Bilenko and Ryan W White. 2008. Mining the search trails of surfing crowds: identifying relevant websites from user activity. In *Proceedings of the 17th international conference on World Wide Web*. ACM, 51–60.
- [4] Christopher M Bishop. 2006. *Pattern recognition and machine learning*. springer.
- [5] Veronika Bogina and Tsvi Kuflik. 2017. Incorporating Dwell Time in Session-Based Recommendations with Recurrent Neural Networks.. In *RecTemp@ RecSys*. 57–59.
- [6] Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger, and Gernot Bauer. 2011. Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage. In *Proceedings of the 13th international conference on Human computer interaction with mobile devices and services*. ACM, 47–56.
- [7] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [8] Christopher JC Burges. 1998. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery* 2, 2 (1998), 121–167.
- [9] Hong Cao and Miao Lin. 2017. Mining smartphone data for app usage prediction and recommendations: A survey. *Pervasive and Mobile Computing* 37 (2017), 1–22.
- [10] Juan Pablo Carrascal and Karen Church. 2015. An in-situ study of mobile app & mobile search interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2739–2748.
- [11] Thomas M Cover, Peter E Hart, et al. 1967. Nearest neighbor pattern classification. *IEEE transactions on information theory* 13, 1 (1967), 21–27.
- [12] David Roxbee Cox. 1966. The statistical analysis of series of events. *Monographs on Applied Probability and Statistics* (1966).
- [13] Ernesto Diaz-Aviles, Hoang Thanh Lam, Fabio Pinelli, Stefano Braghin, Yiannis Kkoufas, Michele Berlingerio, and Francesco Calabrese. 2014. Predicting user engagement in twitter with collaborative ranking. In *Proceedings of the 2014 Recommender Systems Challenge*. ACM, 41.
- [14] Trinh Minh Tri Do and Daniel Gatica-Perez. 2014. Where and what: Using smartphones to predict next locations and applications in daily life. *Pervasive and Mobile Computing* 12 (2014), 79–91.
- [15] Harris Drucker, Corinna Cortes, Lawrence D Jackel, Yann LeCun, and Vladimir Vapnik. 1994. Boosting and other ensemble methods. *Neural Computation* 6, 6 (1994), 1289–1301.
- [16] Claudia Ehmke and Stephanie Wilson. 2007. Identifying web usability problems from eye-tracking data. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI... but not as we know it-Volume 1*. British

Computer Society, 119–128.

- [17] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. Diversity in smartphone usage. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 179–194.
- [18] Denzil Ferreira, Jorge Goncalves, Vassilis Kostakos, Louise Barkhuus, and Anind K Dey. 2014. Contextual experience sampling of mobile application micro-usage. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. ACM, 91–100.
- [19] Steve Fox, Kuldeep Karnawat, Mark Mydland, Susan Dumais, and Thomas White. 2005. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)* 23, 2 (2005), 147–168.
- [20] Andrew Gelman. 2008. Scaling regression inputs by dividing by two standard deviations. *Statistics in medicine* 27, 15 (2008), 2865–2873.
- [21] Google. 2018. *Select a category for your app or game*. <https://support.google.com/googleplay/android-developer/answer/113475?hl=en-GB>
- [22] Balázs Hidasi and Domonkos Tikk. 2016. General factorization framework for context-aware recommendations. *Data Mining and Knowledge Discovery* 30, 2 (2016), 342–371.
- [23] Ryosuke Homma, Keiichi Soejima, Mitsuo Yoshida, and Kyoji Umemura. 2018. Analysis of User Dwell Time on Non-News Pages. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 4333–4338.
- [24] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. 2013. *Applied logistic regression*. Vol. 398. John Wiley & Sons.
- [25] Jeff Huang, Ryen W White, and Susan Dumais. 2011. No clicks, no problem: using cursor movements to understand and improve search. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1225–1234.
- [26] Ke Huang, Chunhui Zhang, Xiaoxiao Ma, and Guanling Chen. 2012. Predicting mobile application usage using contextual information. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 1059–1065.
- [27] Diane Kelly and Nicholas J Belkin. 2004. Display time as implicit feedback: understanding task effects. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 377–384.
- [28] Young Hoon Kim, Dan J Kim, and Kathy Wachter. 2013. A study of mobile user engagement (MoEN): Engagement motivations, perceived value, satisfaction, and continued engagement intention. *Decision Support Systems* 56 (2013), 361–370.
- [29] Farshad Kooti, Mihajlo Grbovic, Luca Maria Aiello, Eric Bax, and Kristina Lerman. 2017. iPhone’s Digital Marketplace: Characterizing the Big Spenders. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 13–21.
- [30] Janette Lehmann, Mounia Lalmas, Elad Yom-Tov, and Georges Dupret. 2012. Models of user engagement. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 164–175.
- [31] Huoran Li and Xuan Lu. 2017. Mining device-specific apps usage patterns from large-scale android users. *arXiv preprint arXiv:1707.09252* (2017).
- [32] Huoran Li, Xuan Lu, Xuanzhe Liu, Tao Xie, Kaigui Bian, Felix Xiaozhu Lin, Qiaozhu Mei, and Feng Feng. 2015. Characterizing smartphone usage patterns from millions of android users. In *Proceedings of the 2015 Internet Measurement Conference*. ACM, 459–472.
- [33] Zhung-Xun Liao, Po-Ruey Lei, Tsu-Jou Shen, Shou-Chung Li, and Wen-Chih Peng. 2012. Mining temporal profiles of mobile applications for usage prediction. In *Data Mining Workshops (ICDMW), 2012 IEEE 12th International Conference on*. IEEE, 890–893.
- [34] Zhung-Xun Liao, Shou-Chung Li, Wen-Chih Peng, S Yu Philip, and Te-Chuan Liu. 2013. On the feature discovery for app usage prediction in smartphones. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*. IEEE, 1127–1132.
- [35] Zhung-Xun Liao, Yi-Chin Pan, Wen-Chih Peng, and Po-Ruey Lei. 2013. On mining mobile apps usage behavior for predicting apps usage in smartphones. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 609–618.
- [36] Akhil Mathur, Nicholas D Lane, and Fahim Kawsar. 2016. Engagement-aware computing: Modelling user engagement from mobile contexts. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 622–633.
- [37] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My phone and me: understanding people’s receptivity to mobile notifications. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. ACM, 1021–1032.
- [38] Cataldo Musto, Giovanni Semeraro, Marco De Gemmis, and Pasquale Lops. 2015. Word Embedding Techniques for Content-based Recommender Systems: An Empirical Evaluation.. In *Recsys posters*.

- [39] Klaas Nelissen, Monique Snoeck, Seppe Vanden Broucke, and Bart Baesens. 2018. Swipe and tell: Using implicit feedback to predict user engagement on tablets. *ACM Transactions on Information Systems (TOIS)* 36, 4 (2018), 1–36.
- [40] Wei Pan, Nadav Aharoni, and Alex Pentland. 2011. Composite social network for predicting mobile apps installation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 25.
- [41] Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 897–908.
- [42] Thanasis Petsas, Antonis Papadogiannakis, Michalis Polychronakis, Evangelos P Markatos, and Thomas Karagiannis. 2017. Measurement, modeling, and analysis of the mobile app ecosystem. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* 2, 2 (2017), 7.
- [43] Martin Pielot, Karen Church, and Rodrigo De Oliveira. 2014. An in-situ study of mobile phone notifications. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. ACM, 233–242.
- [44] Mary Lou Quinlan. 2003. *Just ask a woman: Cracking the code of what women want and how they buy*. John Wiley & Sons.
- [45] Janeaya Revels, Dewi Tojib, and Yelena Tsarenko. 2010. Understanding consumer intention to use mobile services. *Australasian Marketing Journal (AMJ)* 18, 2 (2010), 74–80.
- [46] Jürgen Schmidhuber. 2015. Deep learning in neural networks: An overview. *Neural networks* 61 (2015), 85–117.
- [47] Yoshifumi Seki and Mitsuo Yoshida. 2018. Analysis of user dwell time by category in news application. In *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*. IEEE, 732–735.
- [48] Suranga Seneviratne, Aruna Seneviratne, Prasant Mohapatra, and Anirban Mahanti. 2015. Your installed apps reveal your gender and more! *ACM SIGMOBILE Mobile Computing and Communications Review* 18, 3 (2015), 55–61.
- [49] Choonsung Shin, Jin-Hyuk Hong, and Anind K Dey. 2012. Understanding and prediction of mobile application usage for smart phones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 173–182.
- [50] Iwan Syarif, Ed Zaluska, Adam Prugel-Bennett, and Gary Wills. 2012. Application of bagging, boosting and stacking to intrusion detection. In *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, 593–602.
- [51] Chang Tan, Qi Liu, Enhong Chen, and Hui Xiong. 2012. Prediction for mobile application usage patterns. In *Nokia MDC Workshop*, Vol. 12.
- [52] Maryam Tavakol and Ulf Brefeld. 2014. Factored MDPs for detecting topics of user sessions. In *Proceedings of the 8th ACM Conference on Recommender Systems*. 33–40.
- [53] Yuan Tian, Ke Zhou, Mounia Lalmas, Yiqun Liu, and Dan Pelleg. 2020. Cohort Modeling Based App Category Usage Prediction. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. 248–256.
- [54] Yuan Tian, Ke Zhou, Mounia Lalmas, and Dan Pelleg. 2020. Identifying Tasks from Mobile App Usage Patterns. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2357–2366.
- [55] Steven Van Canneyt, Marc Bron, Andy Haines, and Mounia Lalmas. 2017. Describing Patterns and Disruptions in Large Scale Mobile App Usage Data. In *Proceedings of the 26th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, 1579–1584.
- [56] Henriette C van Vugt, Elly A Konijn, Johan F Hoorn, I Keur, and Anton Eliëns. 2007. Realism is not all! User engagement with task-related interface characters. *Interacting with Computers* 19, 2 (2007), 267–280.
- [57] Theodore Vasiloudis, Hossein Vahabi, Ross Kravitz, and Valery Rashkov. 2017. Predicting session length in media streaming. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 977–980.
- [58] Tianxin Wang, Jingwu Chen, Fuzhen Zhuang, Leyu Lin, Feng Xia, Lihuan Du, and Qing He. [n.d.]. Capturing Attraction Distribution: Sequential Attentive Network for Dwell Time Prediction. ([n. d.]).
- [59] Siqi Wu, Marian-Andrei Rizoiu, and Lexing Xie. 2018. Beyond views: Measuring and predicting engagement in online videos. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.
- [60] Qiang Xu, Jeffrey Erman, Alexandre Gerber, Zhuoqing Mao, Jeffrey Pang, and Shobha Venkataraman. 2011. Identifying diverse usage behaviors of smartphone apps. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 329–344.
- [61] Shijian Xu, Wenzhong Li, Xiao Zhang, Songcheng Gao, Tong Zhan, and Sanglu Lu. 2020. Predicting and Recommending the next Smartphone Apps based on Recurrent Neural Network. *CCF Transactions on Pervasive Computing and Interaction* 2, 4 (2020), 314–328.
- [62] Ye Xu, Mu Lin, Hong Lu, Giuseppe Cardone, Nicholas Lane, Zhenyu Chen, Andrew Campbell, and Tanzeem Choudhury. 2013. Preference, context and communities: a multi-faceted approach to predicting smartphone app usage patterns. In *Proceedings of the 2013 International Symposium on Wearable Computers*. ACM, 69–76.

- [63] Tingxin Yan, David Chu, Deepak Ganesan, Aman Kansal, and Jie Liu. 2012. Fast app launching for mobile devices using predictive user context. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 113–126.
- [64] Xing Yi, Liangjie Hong, Erheng Zhong, Nanthan Nan Liu, and Suju Rajan. 2014. Beyond clicks: dwell time for personalization. In *Proceedings of the 8th ACM Conference on Recommender systems*. ACM, 113–120.
- [65] Elad Yom-Tov, Mounia Lalmas, Ricardo Baeza-Yates, Georges Dupret, Janette Lehmann, and Pinar Donmez. 2013. Measuring inter-site engagement. In *Big Data, 2013 IEEE International Conference on*. IEEE, 228–236.
- [66] Chunhui Zhang, Xiang Ding, Guanling Chen, Ke Huang, Xiaoxiao Ma, and Bo Yan. 2012. Nihao: A predictive smartphone application launcher. In *International Conference on Mobile Computing, Applications, and Services*. Springer, 294–313.
- [67] Sha Zhao, Zhiling Luo, Ziwen Jiang, Haiyan Wang, Feng Xu, Shijian Li, Jianwei Yin, and Gang Pan. 2019. AppUsage2Vec: Modeling smartphone app usage for prediction. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 1322–1333.
- [68] Sha Zhao, Julian Ramos, Jianrong Tao, Ziwen Jiang, Shijian Li, Zhaohui Wu, Gang Pan, and Anind K Dey. 2016. Discovering different kinds of smartphone users through their application usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 498–509.
- [69] Tengfei Zhou, Hui Qian, Zebang Shen, Chao Zhang, Chengwei Wang, Shichen Liu, and Wenwu Ou. 2018. Jump: A joint predictor for user click and dwell time. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. AAAI Press, 3704–3710.
- [70] Xun Zou, Wangsheng Zhang, Shijian Li, and Gang Pan. 2013. Prophet: What app you wish to use next. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication*. ACM, 167–170.