Mobile AD(D)

Estimating Mobile App Session Times for Better Ads

John P. Rula Byungjin Jun Fabián E. Bustamante
Northwestern University
{john.rula, byungjin.jun, fabianb}@eecs.northwestern.edu

ABSTRACT

While mobile advertisement is the dominant source of revenue for mobile apps, the usage patterns of mobile users, and thus their engagement and exposure times, may be in conflict with the effectiveness of current ads. User engagement with apps can range from a few seconds to several minutes, depending on a number of factors such as users' locations, concurrent activities and goals. Despite the wide-range of engagement times, the current format of ad auctions dictates that ads are priced, sold and configured *prior* to actual viewing, regardless of the actual ad exposure time.

We argue that the wealth of easy-to-gather contextual information on mobile devices is sufficient to allow advertisers to make better choices by *effectively predicting exposure time*. We analyze mobile device usage patterns with a detailed two-week long user study of 37 users in the US and South Korea. After characterizing application session times, we use factor analysis to derive a simple predictive model and show it is able to offer improved accuracy compared to mean session time over 90% of the time. We make the case for including predicted ad exposure duration in the price of mobile advertisements and posit that such information could significantly impact the effectiveness of mobile ads by giving publishers the ability to tune campaigns for engagement length, and enable a more efficient market for ad impressions while lowering network utilization and device power consumption.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Communications Applications

General Terms

Experimentation; Measurement; Performance

Keywords

Mobile; Apps; Ads; Prediction

HotMobile'15, February 12–13, 2015, Santa Fe, New Mexico, USA.

ACM 978-1-4503-3391-7/15/02 ...$15.00.
http://dx.doi.org/10.1145/2699343.2699365.

1. INTRODUCTION

Advertisement is the dominant source of revenue for mobile apps, with nearly 90% of available apps offered for free. Mobile ad sales more than doubled between 2012 and 2013, totaling over $17.96 billion, and are projected to rise another 75% in 2014 alone [5]. Despite this impressive growth, we posit that the effectiveness of current ad campaigns may be hindered by the usage patterns of mobile users, and their engagement times.

Ad campaigns are aimed, among other goals, at improving online site traffic, creating advertising recall, brand recognition [1] and brand awareness [4]. Display ad campaigns target exposure, rather than site traffic or sales, and have traditionally used a pricing scheme based on the number of impressions delivered (measured as CPM or cost per thousand impressions). Following a model inherited from the newspaper and a mostly-static Web era [6], online display ads treat all impressions the same, independently of the total exposure time, despite the clear benefits that longer exposure has on recognition and recall [6, 7].

While user engagement with mobile apps can range widely, from seconds to several minutes, depending on factors such as users' locations, concurrent activities (e.g., running, sitting on a train) and goals (e.g., entertainment, finding direction, work), we argue that the usage patterns and wealth of easy-to-gather contextual information on mobile devices is sufficient to *effectively predicting session or exposure time*.

In mobile settings, where there is only one app in the foreground, ad exposure time is bound by application session time. Knowing either can benefit all parties in the advertisement ecosystem. Advertisers can use this information to tune campaigns for engagement length, and bid on the appropriate value of an impression. Session time information could reward publishers for engaging users, give ad networks additional freedom to optimize their selections, and reduce wasted resources on end host devices by eliminating data for ads which are never shown. Surprisingly considering its many benefits, we are not aware of any study to attempt to predict the length of a mobile application session or ad exposure times through contextual factors.

We make the case for including predicted ad exposure time in the price of mobile advertisements in current ad exchanges where impressions are auctioned at their onset. Using a detailed two-week-long user study, we analyze the mobile device usage behavior of 37 users (200,000 application sessions) in two mobile markets – the US and South Korea. We show that application session times form a long tailed power law distribution, implying a large disparity in the quality and value of mobile ad sessions. We employ factor

[1] *Recall* is the proportion of users who report remembering an advertising with a minimal prompt, while *recognition* uses text or images as probes.

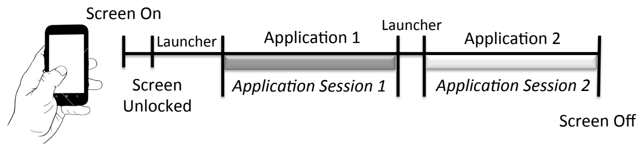

Figure 1: Diagram of actions recorded by our measurement service. We are only interested in the time a particular app is both visible and in the foreground on a user's device, indicated as *Application Session 1* and *Application Session 2* in the Figure.

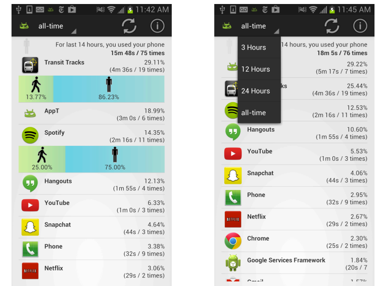

Figure 2: Two snapshots of *AppT* illustrating a subset of the applications monitored, their usage time, and activity, as well as the observation intervals supported (3, 12 and 24 hours and all-time).

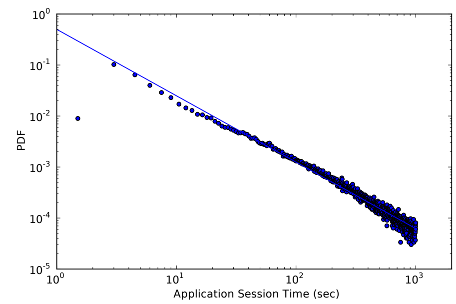

Figure 3: Distribution of application session times form a power law distribution in the form of $f(x) = 0.499x^{-1.3014}$. Average application session time from our dataset is 258 seconds.

analysis of device contextual components to determine dominant influences in device usage and application session time and to inform the design of our predictive models. We then show that our prediction model improves accuracy over mean session time over 90% of the time. Our preliminary results show that even a simple first-approximation model to predict session time can significantly improve over current practices of ad campaigns.

2. BACKGROUND AND MOTIVATION

There are a number of common billing models for online advertisement, including Cost Per Click (CPC), Cost Per Mille or thousand impressions (CPM), and Cost Per Action or Acquisition (CPA). Independently of the billing model, advertisers typically purchase ad space by bidding in an auction format.

Bids for impressions are based on the user target profile, which includes demographics such as gender, age and purchasing power, and mobile application category. For instance, certain (classes of) users are particularly coveted due to factors such as their interest in the advertisers' subject or their purchasing power. Ads on more popular websites or applications are also worth more, as are advertisements in more prominent locations on websites.

To the best of our knowledge *no ad auction today takes into account the (expected) time a user is engaged with an application (or advertisement)*, despite the known benefits of longer session times [6]. A simple approach to incorporate session time would be to rely on average impression times during bids. As we show in the following section, considering the high variance and expected long tail distributions of app session times, assuming average values would not be particularly useful.

We argue that, unlike their traditional online counterparts, mobile ads are well suited to include accurate temporal information in their advertisements. The constraints of mobile devices allow mobile applications to accurately measure the amount of time a user is exposed to an advertisement. For starters, in mobile setting, there is only one application that can be in the foreground at a time. This eliminates the ambiguity of having multiple windows, or browser tabs open simultaneously. Device usage is bounded by the time the screen is illuminated, and user interaction can be ensured by the progression through user identification mechanisms such as lock-screens. This ensures that a user is actively engaged with the application within a fine margin. An illustrative example of mobile device and application use flow is shown in Figure 1.

Knowing the session length of a publisher's impression can benefit all parties in the advertising ecosystem, from advertisers themselves to end users. This information could allow advertisers to actually pay for the amount of time their ads are shown. In addition, advertisers can tune their ad campaign media for the appropriate session length. Similarly, publishers with long user engagement can be accordingly compensated. Knowing the session length during an ad request allows the ad network the flexibility to adjust the ordering of ads, and multiplex ad sessions if desired. Last but not least, knowing session time, allows advertising libraries to avoid wasting network resources and cap bandwidth by only downloading the ads which are needed for the allotted time [10].

While there exist many benefits to this type of mobile advertising model, however, if would require a reinvention of the existing advertising marketplace. We instead propose that the same benefits can be enjoyed by predicting the length of a mobile application session, and using this information in existing online ad auctions. Our idea leverages the rich context available mobile devices to build predictive models of application session times. Mobile devices have access to a wealth of contextual information for an app usage, including the user's location, activity, historical usage patterns, and network performance – context which is unavailable for desktop users. Given that error would be inevitable in these predictions, we formalize the cost of prediction error and calculate the cost of this error in the following section.

3. COLLECTING APP SESSION TIME

To explore the possibility of predicting app session time and determine the set of contextual factors necessary for this, we conducted a field study tracking the usage patterns of a set of real users. To this end, we developed and made available an application we call *AppT* – for *Application Time*. AppT tracks application session times, defined as the length of time a mobile app is in the foreground on the user's device (illustrated previously in Fig. 1). Figure 2 presents two screenshots of our app.

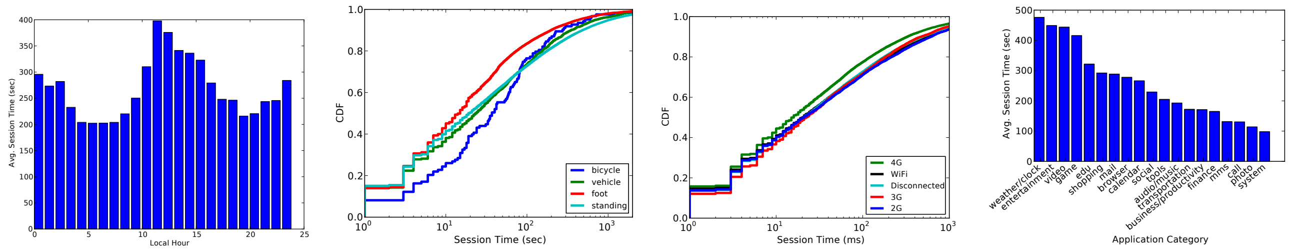

(a) By hour. (b) By activity. (c) By network connectivity. (d) By app category.

Figure 4: Application session times for users under different contexts.

We consider each application to be the *foreground application* if it visible to the user (i.e. the screen is illuminated and not behind a lock screen), and the first ranked foreground application by the Android operating system. For the duration of screen illumination, we polled the foreground application from the operating system every 1.5 seconds. In addition to application session time, AppT records contextual information from each participant's device, including the screen illumination time, network conditions, and user activity.

For this study, we use AppT collected detailed application usage time from 35 users, in the United States and South Korea, during a two week periods in March and April 2014. Our dataset includes over 200,000 individual application sessions.

3.1 Mobile App Usage

We find application session times to follow a power law distribution. Figure 3 plots the distribution of application session time, showing that it forms a power law in the form of $f(x) = 0.499x^{-1.3014}$. A long tail distribution means a highly skewed mean session time. For instance, the average session time from our dataset was 258 seconds, however, this value represents the 90th percentile of the entire distribution.

4. CONTEXTUAL FACTOR ANALYSIS

We perform, to the best of our knowledge, the first analysis of contextual factors on mobile application session time. Using the collected dataset we explore the impact of user activity, network connectivity and performance, and temporal components on mobile application session time, and use this analysis to inform the design of our prediction models. We expect that different (types of) applications will be use and be impacted differently by the different factors we explore. Rather than assuming a model per application, we explore the use of an application name/class in our contextual analysis.

We want to identify those that are most dominant and informative for predicting application session time. To this end, we use two quantitative approaches for comparing factors: analysis of variance (ANOVA) tests for statistical independence between categories within a factor (e.g., standing, walking, bicycling, in a vehicle) and information gain analysis to measure the decrease in entropy per category.

In our context, high statistical independence between categories hint at the relative value of that contextual factor as a predictor of application session time. Similarly, high information gain (i.e. entropy reduction) indicates greater predictability of application session time based on the factor's categories. As an illustrative example, if one would like to predict the height of an Olympic athlete, a high statistical independence between each athlete's sport (e.g. between basketball and figure skating), of the height distribution grouped by sport, would indicate that *sport*, can serve as a good predictor.

4.1 Contextual Factors

The rich information available on the context of a mobile device usage has proved valuable to make application launch prediction [12, 15, 16]. Our hypothesis is that this information can be further leverage to accurately predict application session time.

A description of the contextual factors we gathered in our study are given below.

Temporal Components. We expect temporal information, such as time of day or day of the week, to have an effect on the total usage and application session length on users' devices. To view temporal trends in application session time, we look at how app session lengths change over time. We analyze how session times change with regard to the *hour of the day* and the *day of the week*.

Figure 4a shows the average session time for our study users binned by each hour of the day. The figure shows bimodal usage peaks each day, with application session length peaking at 2 pm and 12am local time. In addition, we observed larger peaks of high session times on Saturdays and Sundays than during weekdays.

User Activity. Given the level of integration of mobile devices into our everyday lives, we also expect user activity to have a dominant role in determining application session time. For instance, an exercise tracking application might have a much higher probability of being used while the user is running or cycling, however, the total engagement time during each session might be much shorter during the activity than later, when they are reviewing their performance.

To this end, AppT records the *current user activity* as taken from the DetectedActivity intent built into Android operating system. Activities are detected in 20 second granularities, recording whether the user is walking, standing, in a vehicle, running or bicycling.

Figure 4b plots session length per activities for all users. In aggregate, walking and cycling sessions show, respectively, the shortest and longest median session times with respective values of 16 and 41 seconds. Interestingly, both stationary and vehicular sessions are similarly distributed but present significantly different average session times (262 and 175 seconds).

Network Connectivity Besides the time and current activity of a user when in an application, one would expect the quality and performance of the device's network connection to impact application session times. For instance, a messaging application such as WhatsApp would be of little use without any network connectivity, while poor network conditions tend to render a mobile web browser virtually unusable. We captured each device's current *connectivity state* (connected or disconnected) along with the current *radio interface* (e.g. WIFI, LTE, UTMS, etc) in use during a connected period.

We found individual network connectivity states to offer very similar distributions for application session time. Figure 4c plots the distribution of application session times under different connectivity conditions (e.g. 4G, 3G, WiFi, etc). The figure shows almost identical curves for all connectivity states with the exception of LTE, which shows shorter application session than each other network state.

Mobile Application Given the wide range of applications available and the different intended usage modes, we expect that either the application or its type to be key for predicting session time prediction. For our analysis, we use the application's *package name*, which uniquely identifies it within the Android ecosystem. We focus our study on the top ten most commonly used applications for users. We found this subset to be sufficient to account for most of the device utilization – indeed, we found the top 5 applications alone already account for over 80% of device usage, on average.

Application sessions also differ based on the category of application. We categorized each application package into one of 19 different categories (e.g. messaging, game, video). We found the category of mobile application to play a significant determinant of average session time. Average session times for each category are plotted in descending order in Figure 4d. The figure shows, surprisingly, that common phone utilities such as weather/clock had the highest average session times of nearly 3 minutes. Other apps such as messaging and phone contacts had some of the lowest average session times.

4.2 Analysis of Variance (ANOVA)

ANOVA tests are used to determine independence between subpopulations of a population (categories of a factor, in our context). ANOVA represents independence between categories by looking at the overlap between confidence intervals of each class. The formula shown by Equation 1 calculates the F ratio, which is then plugged into the standard F-distribution to obtain p-values for significance [8].

$$F = \frac{MS_{Treatments}}{MS_{Error}} = \frac{\frac{\sum n_j(\bar{x}_j - \bar{x})^2}{k-1}}{\frac{\sum (n_j - 1)\sigma_j^2}{n-k}} \quad (1)$$

We ran ANOVA tests, per user, for each factor described previously in Section 4.1. The significance levels are represented as p-values for each factor and shown in Figure 5. Typically p-values less than 0.05 (dash vertical line) indicate high levels of subpopulation independence, while p-values less than 0.1 (solid vertical line) indicate weak independence [8].

We use these results to determine which factors should be included within our prediction model. Those with high group independence (low p-values) mean that grouping by that category produces statistically significant differences in each subpopulation, indicating a factor will be useful for prediction.

Our results (Fig. 5) show hour of the day to have the highest group independence, followed by mobile application, user activity and day of week. Indeed, time of day shows high levels of group independence (p-values < 0.05) for over 90% of users. Mobile application showed similar independence also for over 90% of users. User activity obtained the next highest , with weak independence for nearly half of the user population, and radio type seeing independence in nearly 30% of the user population. Connectivity was the least independent of all factors, showing strong independence in only 10% of the user population and weak independence in only 20% of users. Unlike time of day (hour),

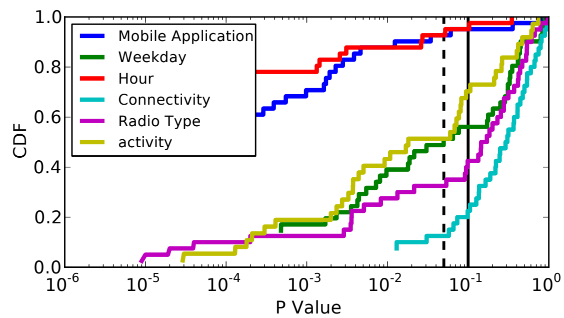

Figure 5: P-values for users for different contextual groups. P-values < 0.05 (dotted line) indicate statistical significant independence, while p-values < 0.1 (solid line) indicate weak independence.

each of these last factors represent minor indicators of application session time according to statistical independence. We supplement these insights with information gain analysis in the next section.

4.3 Information Gain Analysis

To further characterize contextual factor influence on application time, we calculate the *relative information gain* of each factor by measuring the decrease in entropy each brings to the system, or more plainly, the increase in predictability of session times each category of contextual factor gives. The information gain of each contextual factor can be used to guide the design of predictive models.

The entropy of a random variable Y is given as $H(Y) = \sum_i P[Y = y_i] log \frac{1}{P[Y=y_i]}$, where $P[Y = y_i]$ is the probability that $Y = y_i$. The information gain $G(X)$ of a particular factor F with states $f \in F$ is calculated as $G(Y, F) = H(Y) - \sum_{f \in F} \frac{|Y_f|}{|Y|} H(Y_f)$. It is the difference between the total entropy of the original system, and the sum of the entropy of each factor grouping. For comparison purposes, we calculate the *relative* information gain, which normalized the difference in entropy by dividing the result by the total entropy of the system, $H(Y)$.

Figure 6 plots the cumulative distribution of information gain observed by each user in our study for the 5 contextual factors. The figure shows, again, that the temporal components (weekday and hour) along with application name (package name) provide the highest amounts of information gain to the system. Radio type offered moderate information gains, while user activity gave the least by far.

Interestingly, user activity offers the lowest information gain even though our ANOVA analysis (§ 4.2) showed there exists high level of independence between activity subpopulations. This is due to the fact that information gain is not normalized against subpopulation size like ANOVA, and speaks to the need for multiple techniques when designing a prediction model. In the case of user activity, the number of application sessions classified as *stationary* were at least an order of magnitude larger than any of the other activities. Therefore, the entropy of *stationary* sessions (which is weighted by $\frac{|Y_f|}{|Y|}$) is very close to the total entropy of the system.

5. MOBILE APP PREDICTION

We now outline a procedure for modelling and predicting application session times. We first formalize the bounds of prediction error by calculating the maximum error allowable to still offer

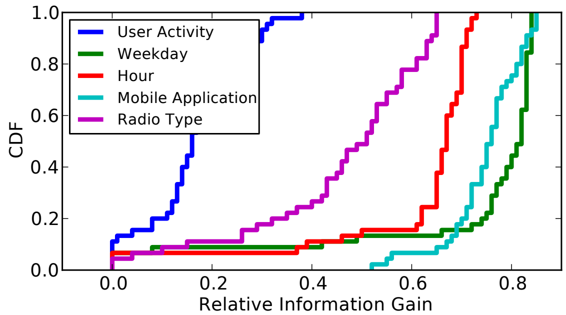

Figure 6: Relative information gain for contextual factors for application session time. Application name, and temporal components have large information gains for study users.

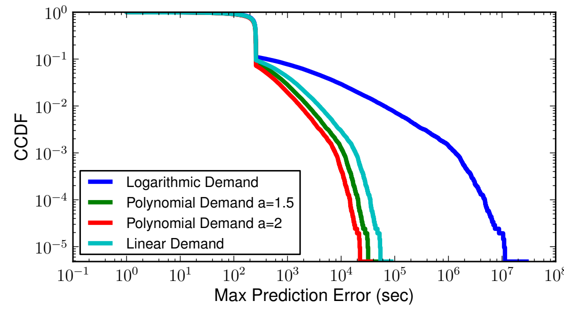

Figure 7: Maximum prediction error ($e_t = |t_p - t|$) tolerable for session prediction to improve against the current model. While each demand model differs on its tail behavior, errors for each demand model are very similar, and heavily dependent on $\bar{t}$, which in our study is 258.9 seconds.

more accurate information than the average session time. Using this model, we compute the potential cost of inaccurate prediction on advertisers and publishers before presenting preliminary work toward session time prediction.

5.1 Prediction Error Bounds

Since ad sales are closed *before* the user has even begun viewing an ad, it is impossible to know the actual session length at bidding time, making it necessary to use an estimate. We formalize the bounds of prediction error by calculating the maximum error allowable to still offer more accurate information than the average session time, $\bar{t}$. With prediction, there always exists some error, e_t, between the predicted session time, t_p, and the actual session time, t.

We define the loss due to prediction error as the difference in price between the predicted session time, and the price of the actual session time, $|P(t_p) - P(t)|$. This cost due to any prediction error can only be determined by knowing the shape of the advertiser demand function, $P(t)$. This is due to our definition of loss being based on the price differential between actual and predicted session times, therefore the shape of the price (demand) curve is integral to the overall loss. We compare this prediction loss to the difference in price between that of mean session time and the price of the actual session time, $|P(\bar{t}) - P(t)|$.

We look at the results from our estimation of advertiser demand curves using linear, polynomial and logarithmic growth functions. We look for conditions where the loss from prediction is less than the loss from the current model.

$$|P(\bar{t}) - P(t)| > |P(t_p) - P(t)| \tag{2}$$

Using linear demand where $P(t) = Ct + b$, we can reduce Equation 2 to $|\bar{t} - t| > |t_p - t|$, or more simply, when the error from prediction is less than the distance to the session time mean. Using a polynomial demand where $P(t) = Ct^a + b$ and $a > 1$, we can reduce Equation 2 to $|\bar{t}^a - t^a| > |t_p^a - t^a|$, the value of session time here depends on the magnitude of the demand exponent, along with prediction accuracy. Using logarithmic demand where $P(t) = Clog(t) + b$, Equation 2 reduces to $|log(\frac{\bar{t}}{t})| > |log(\frac{t_p}{t})|$.

Using the distribution of application session times collected from our user study (§ 3), we simulate the maximum value of prediction error, e_t which can be tolerated for our prediction model using the three possible demand curves described above. Figure 7 shows the maximum prediction error for Equation 2 to hold using all application session times taken from our experiments. The Figure shows the high similarity in allowable prediction error for the different demand curves, diverging only above the 90th percentile among all curves. This similarity shows that for the vast majority of cases, the advertiser demand function has a minimal effect on the bounds of prediction error.

5.2 A Naïve Prediction Model

Using the data collected from our experiments (§ 3), and our contextual factor analysis (§ 4) we constructed and evaluated a naïve prediction model for mobile application sessions. Due to the complex interactions between contextual factors and app session time, we chose to use a decision tree classifier to predict session times over other classifiers and regression models. Each class is generated taking an equal percentile from the overall training set distribution. Using a classifier over potential regression models reduces overall accuracy since the predicted value is the average value taken for a given classification; however, we found this to be more accurate when compared to regression models such as linear or decision-tree regression.

We analyzed the accuracy of our predictive models through the classification accuracy, the percentage of session times which are correctly placed in the right class. To evaluate our prediction model, we split our dataset into a training and validation set. Since our dataset encompassed a two week period, we use the first week of data for model training and the second week as a validation set for our models. This allows us to compare the success of several different classifier models and class sizes. We found that decision trees provided the most accurate predictor of mobile application session times, when compared with other common classifiers such as support vector classifiers (SVCs) and Naive Bayes classifiers. Unsurprisingly we found that as we increased the number of classes we see classifier accuracy decrease substantially.

Application session times are continuous values, and the classifier accuracy does not capture the absolute error obtained by the prediction. For instance, if a session was classified incorrectly, but still placed in an adjacent class, the prediction might still be beneficial. We therefore also measure the absolute prediction error of our predictive model, defined as the difference between the predicted session time and the actual session time from our testing set. Since our predictions are based on classes, the predicted time is taken to be the average session time in each class from the training set. Figure 8 shows this prediction error for our prediction model, the mean error, and a random distribution sampling. Our classifier outperforms mean error in over 90% of cases.

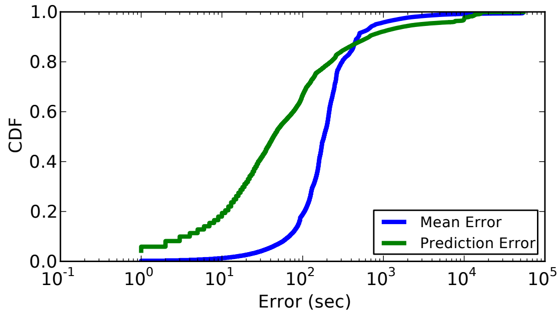

Figure 8: Prediction error for simple decision tree classifier and 10 classes, compared to the mean error. Our prediction model outperforms mean error in over 90% of cases.

6. RELATED WORK

Several recent efforts have looked at the impact of device context on mobile device usage patterns, and understanding user engagement and application usage for individuals. Research on user engagement has focused on determining dominant factors for user viewing or for developing predictive models for engagement in online videos (e.g., [1,3]) .

Other projects have explored application usage and prediction on mobile devices (e.g., [2, 12, 13, 15, 16]) considering contextual factors such as time of day and location to predict, for instance, the next application to be launched. Understanding the effect of context on device usage has important implications on system performance enhancements, device preloading and prefetching, and application design.

Our work is complementary to these efforts. We approach the problem of usage prediction not from the individual applications, but from the level of the entire device, and bringing in additional contextual information such as user activity (e.g., walking, standing), network state, total phone session and performance. By focusing on device engagement, we are able to understand usage patterns across classes of applications and over short individual application sessions. Ourgoal is to better understand the effect that each of these contextual elements has on device and application usage, and, use this information to reliably predict session times for devices and applications.

Due to their large role in financing mobile applications, mobile advertisements have recently been studied in the contexts of fraud [9], contextual effectiveness [11], and network and power usage on mobile devices [14]. Closest to our work is by Mohan et al. which studied the efficacy of prefetching mobile advertisements [10]. Part of their analysis was to look at the entropy of application session lengths from a large historical sample, finding that prediction is indeed feasible. Our work extends their initial analysis, taking into account a large variety of user contexts, and implementing a first attempt at mobile session prediction from this context.

7. CONCLUSION AND FUTURE WORK

The usage patterns of mobile users, and thus their engagement times, may be in conflict with the effectiveness of current ads. While users engagement with mobile apps can range from a few seconds to several minutes, the current format of ad auctions dictates that ads are priced, sold and configured prior to actual viewing, *regardless* of the actual exposure time. We argue that the wealth of easy-to-gather contextual information on mobile devices is sufficient to allow advertisers to make better choices by *effectively predicting exposure time*. Building on a two-week-long user study in two markets we analyzed mobile device usage patters. We used factor analysis to derive a simple predictive model and show that is able offer improved accuracy compared to mean session time over 90% of the time. We made the case for including predicted ad exposure time in the price of mobile advertisements and posit that such information could significantly impact the effectiveness of mobile advertisement, giving publishers the ability to tune campaigns for engagement length and enabling a more efficient market for ad impressions, select appropriate media for an ad impression and lowering the cost to users including network utilization and device power. In ongoing work, we are exploring better prediction models and evaluating the benefits of estimated session times in other contexts, including media selection and network usage.

8. ACKNOWLEDGEMENTS

We would like to thank our shepherd, Alex Snoeren, and the anonymous reviewers for their valuable feedback and assistance. This work was supported in part by the National Science Foundation through award CNS 1218287.

9. REFERENCES

[1] A. Balachandran, V. Sekar, A. Akella, S. Seshan, I. Stoica, and H. Zhang. Developing a predictive model of quality of experience for Internet video. In *Proc. ACM SIGCOMM*, 2013.

[2] M. Böhmer, B. Hecht, J. Schöning, A. Antonio Krüger, and G. Bauer. Falling asleep with angry birds, facebook and kindle: A large scale study on mobile application usage. In *Proc. of MobileHCI*, 2011.

[3] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang. Understanding the impact of video quality on user engagement. In *Proc. ACM SIGCOMM*, 2011.

[4] X. Dreze and F.-X. Hussherr. Internet advertising: Is anybody watching? *Journal of interactive marketing*, 17(4):8–23, 2003.

[5] eMarketer. Driven by facebook and google, mobile ad market soars 105% in 2013 - emarketer. `http://bit.ly/1lOsFlh`.

[6] D. G. Goldstein, R. P. McAfee, and S. Suri. The effects of exposure time on memory of display advertisements. In *Proc. of EC*, 2011.

[7] D. G. Goldstein, R. P. McAfee, and S. Suri. Improving the effectiveness of time-based display advertising. In *Proc. of EC*, 2012.

[8] M. H. Kutner, C. J. Nachtsheim, and J. Neter. *Applied Linear Regression Models*. McGraw-Hill/Irwin, fourth international edition, 2004.

[9] B. Liu, S. Nath, R. Govindan, and J. Liu. Decaf: detecting and characterizing ad fraud in mobile apps. In *Proc. USENIX NSDI*, 2014.

[10] P. Mohan, S. Nath, and O. Riva. Prefetching mobile ads: Can advertising systems afford it? In *Proc. of Eurosys*, 2013.

[11] S. Nath, F. X. Lin, L. Ravindranath, and J. Padhye. Smartads: bringing contextual ads to mobile apps. In *Proc. of MobiSys*, 2013.

[12] A. Parate, M. Böhmer, D. Chu, D. Ganesan, and B. M. Marlin. Practical prediction and prefetch for faster access to applications on mobile phones. In *Proc. of UbiComp*, 2013.

[13] C. Shin, J.-H. Hong, and A. K. Dey. Understanding and prediction of mobile application usage for smart phones. In *Proc. of UbiComp*, 2012.

[14] N. Vallina-Rodriguez, J. Shah, A. Finamore, Y. Grunenberger, K. Papagiannaki, H. Haddadi, and J. Crowcroft. Breaking for commercials: characterizing mobile advertising. In *Proc. IMC*, 2012.

[15] Y. Xu, M. Lin, H. Lu, G. Cardone, N. Lane, Z. Chen, A. Campbell, and T. Choudhury. Preference, context and communities: a multi-faceted approach to predicting smartphone app usage patterns. In *Proc. of ISWC*, 2013.

[16] T. Yan, D. Chu, D. Ganesan, A. Kansal, and J. Liu. Fast app launching for mobile devices using predictive user context. In *Proc. of MobiSys*, 2012.