Understanding Short-term Changes in Online Activity Sessions

Farshad Kooti
Facebook Inc.
Menlo Park, CA

Karthik Subbian
Facebook Inc.
Menlo Park, CA

Winter Mason
Facebook Inc.
Washington, D.C.

Lada Adamic
Facebook Inc.
Menlo Park, CA

Kristina Lerman
USC ISI
Marina del Rey, CA

ABSTRACT

Online activity is characterized by regularities such as diurnal and weekly patterns, reflecting human circadian rhythms and work and leisure schedules. Using data from the online social networking site Facebook, we uncover temporal patterns at a much smaller time scale: within individual sessions. Longer sessions have different characteristics than shorter ones, and this distinction is already visible in the first minute of a person's session activity. This allows us to predict the ultimate length of his or her session and how much content the person will see. The length of the session and other factors are in turn predictive of when the individual will return. Within a session, the amount of time a person spends on different kinds of content depends on both the person's demographic attributes, such as age and the number of Facebook friends, and the length of the time elapsed since the start of the session. We also find that liking and commenting is very non-uniformly distributed between sessions. Predictions of session duration and activity can potentially be leveraged to more efficiently cache content, especially to mobile devices in places with poor communications infrastructure, in order to improve user online experience.

Categories and Subject Descriptors

H.4.3 [**Information Systems**]: Information systems applications

Keywords

Information consumption; activity session; Facebook; prediction

1. INTRODUCTION

Online activity exhibits strong temporal regularities on daily, weekly, and seasonal scales. These regularities have been observed across an array of platforms: voting for news stories submitted to the social news aggregator Digg displays clear daily and weekly cycles of activity [25]; moods expressed by Twitter users worldwide show daily and seasonal variation [5]; daily patterns of food consumption, as well as increased nightlife activity on the weekends, emerge from Foursquare check-in data [9]. These patterns can be attributed to activity states governed by circadian rhythms, sleep cycles, seasonal changes in day length, and work and leisure schedules. In this paper, we uncover regular changes in online activity that take place on an even shorter time scale: minutes, instead of hours or days. Understanding these changes will help predict individuals' behavior and potentially allow for improving their online experience, for example, by informing caching algorithms that anticipate individuals' information needs.

We conduct our study using data from the popular social networking service Facebook, which is used daily by more than a billion people worldwide to stay in touch with family and friends, to connect with communities and interests, to be informed about current events, and to be entertained. Like many other social networking services, Facebook compiles stories shared by friends, pages, and groups, which includes status updates, photos, videos, links to other online content, etc., and presents it as a list (i.e., News Feed). A person browses this list to read status updates from friends or watch photos and videos they shared.

One of the challenges of working with large-scale observational data is that human behavior is highly heterogeneous. For Facebook users, this translates into large variation in preferences about how they read the News Feed (on a mobile device or web browser), how much of the News Feed they read, and so on. To partly control for this heterogeneity, we segment the time series of an individual's activity on Facebook into *sessions*. We define a session as a series of consecutive interactions without a break longer than 10 minutes. By comparing sessions of the same length, we find that individuals spend less time viewing each story as the session progresses. In addition, we find that people preferentially shift attention to some types of content, such as photos, over the course of a session. These trends are more pronounced in the older population and also in people who have fewer friends on Facebook. Moreover, longer sessions have markedly different patterns of activity than shorter sessions. This distinction is so strong that we can use the first minute of a person's activity to predict how long his or her session will be. We can also predict how much content the person will consume over the course of a session, and when he or she will return to Facebook.

The main contributions of this paper are as follows:

- We demonstrate short-term changes in activity, with people spending less time on each story in the News Feed over the

WWW'17 Companion, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4914-7/17/04.
http://dx.doi.org/10.1145/3041021.3054203

course of a session. The rate of these changes varies for different demographic segments.

- We show that as the session progresses, people change their patterns of content consumption, e.g., they spend more time viewing photos rather than textual posts.
- We predict the length of an individual's session using only the activity during the first minute of the session, more accurately than a competitive baseline. We also predict how many stories an individual will consume and the time he or she will return to Facebook.
- We characterize some of the variety of session types, including sessions where people are more likely to comment on stories, sessions where they prefer to "like" stories, and sessions where they mostly read the News Feed.

Although our work does not resolve the origins of these behavioral changes, quantifying them and moreover, using them to predict behavior, can potentially allow for improvements in user online experience. For example, content could be ranked dynamically to account for these behavioral changes by shifting photos to later in the session. The ability to predict session length and activity could be particularly useful for caching content on mobile devices. In developing countries and emerging markets, there are hundreds of millions of users with out-dated mobile devices and poor internet connectivity. Correctly caching content on such devices, based on the predicted session activity—and at the right time (just before the user logs-in, based on predicted return times)—can potentially improve the overall user experience by minimizing the network latency of delivering fresh content.

2. RELATED WORK

Multiple studies have shown daily, weekly, monthly, and yearly patterns of activity in offline and online world. Grinberg et al. [9] show daily and weekly patterns of eating, drinking, shopping, and nightlife in human behavior using Foursquare checkins. Golder et al. [6] found consistent weekly and seasonal patterns of social interaction among college students on Facebook. Later, Golder and Macy [5] drew a connection between sentiment on Twitter posts to cycles of sleep and seasonality. Naaman et al. [15] studied the variations of keyword use on Twitter diurnal patterns and assessed their robustness across geographical locations. Leskovec et al. developed a framework for tracking variants of short textual phrases over time [15] and found prototypical temporal patterns in the spread of news stories [27]. Moreover, it has been shown that people tend to reply to emails faster during the mornings and in weekdays [13], that email activity is bursty yet predictable [16] and that even postal letter replies exhibit certain regularities [17].

Sessions of activity have been constructed to understand online behavior especially in the context of search and web surfing [22, 19, 3]. In different studies sessions could refer to a set of actions to satisfy a single information need [4, 11], or more commonly, a period of time that includes consecutive actions without a long break [24, 7]. In recent work, Kapoor et al. proposed a hidden semi-Markov model to predict the song to recommend to an user, based on their sessionized music listening history [12].

In recent years, sessions of activity have been studied in online social networks, as well. Benevenuto et al. have aggregated data from multiple social networks and created sessions of activity to understand high-level behavior of users in usage of online social networks, e.g., how frequently people login to these sites and how much time they spend browsing them [1, 2]. Focusing on Facebook sessions, Grinberg et al. looked at the effect of contribution, i.e. when an individual posts content of their own on Facebook, and how this correlates with the number and length of sessions, as well as interaction with different types of content [8].

Caching is one area where being able to predict session start and duration is important. The use of mobile devices has increased consistently in the recent years, but there are still areas with poor network connection. Prefetching is one of the solutions to mitigate the problem by downloading the content that the person will need before she or he requests it. Shoukry et al. [21, 20] have implemented this idea by leveraging the predictability of people's behavior to prefetch content while the person is on a wireless network. Another recent study proposed a content pre-fetcher for Twitter that uses signals in the social data to retrieve news feeds and links ahead of usage [26]. Features used included user and tweet characteristics, and the system was able to estimate the users' content interest fairly accurately, resulting in reducing the delay. A pattern-based approach has been used to pre-fetch applications in the memory of mobile devices to decrease application startup time [23]. Note that none of these works have focused on understanding the online user behavior, particularly the length and break-time prediction, in order to inform the pre-fetching schedule.

3. METHODS

In this section, we explain our data collection methodology and data processing steps. All analyses were performed in aggregate on de-identified data, where the users have been anonymized.

3.1 Data

Facebook is an online social network that is used daily by more than a billion people worldwide to stay in touch with family and friends, to connect with communities and interests, to be entertained, or to be informed about current events. A primary activity on Facebook is browsing the News Feed to consume *stories* shared by friends, which include friends' status updates that can be in form of textual posts, videos or photos they shared. By default, the News Feed ranks all the friends' stories by their predicted relevance and interest to the user. Since we are interested in the short-term behavioral changes, such as those occurring over the course of a session, the News Feed ranking algorithm may introduce a substantial confounding factor, for example, by putting more interesting stories higher in the News Feed, so that a person will see them earlier in a session.

Facebook also allows users to rank the stories in chronological order, with the most recently shared story at the top of the list. This option is called "most recent" ranking, and although just a small fraction of people use it, they represent a large enough sample to test our hypotheses.

For our study, we considered only the people who chose the "most recent" option for ranking stories in their News Feed. As a result, the stories they saw on Facebook were ordered by the time of story posting, rather than relevance, so that any observed differences in engagement with stories would be due to factors such as time spent in the session, rather than changes in properties of the stories. Also, since most of the users log in to Facebook multiple times a day, there is not a significant difference in the time that the stories have happened. On average, the population using most recent ranking is more active than the general Facebook population, but they are broadly distributed across different demographic segments such as age, location, and number of friends.

We consider all activities for a random sample of these people over the course of one month (June 2015). In addition, we focus only on people who used Facebook via the web, an iOS device, or an Android device. This sample of millions of users performed billions of interactions on each of the platforms considered. Here, an action is any type of activity that a person can perform on Facebook, such as reading a story, liking or commenting on it, or creating a status update (including links, pictures, text, etc.). In Section 4, we consider a random day of user interactions for our analysis as the results are consistent across all days in the observation period.

We first consider the consumption of information, comprised of: 1) reading a story, which includes reading textual stories and comments, and viewing photos, and 2) watching videos, which includes watching both videos that are originally shared on Facebook and videos that are shared from external websites.

Next, we calculate the time people spent on different activities, such as reading posts and watching videos in their News Feed. We used the logged data to calculate the time that a person spends on each story. This is achieved by considering all stories that are visible to the user in the News Feed, and dividing the time spent viewing between the stories based on the proportion of the screen that the stories are occupying. For example, if there are only three stories visible in the screen and the first one occupies half of the screen, the second one 20%, and the third one 30% of the screen, and the user spends 10 seconds on reading these stories, then we give an approximate allocation of 5 seconds spent on the first story, 2 seconds on the second story, and 3 seconds on the last story.

For video we use the amount of time a video has been played in the News Feed. By default, videos in the News Feed are played automatically, and we count them as watched if the viewer switches to full-screen, un-mutes the sound, or stays on the video for at least 75% of its length. Other thresholds and criteria yielded similar results.

We study how people allocate their time to read stories and watch videos and how this allocation changes over the course of an activity session.

3.2 Activity Sessions

To study changes in behavior over the course of a session, we have to segment the time series of user interactions data to identify sessions. One option is to use the actual sessions; that is, the time beginning when a person navigates to Facebook or opens the Facebook application until the time the web page or application is closed. However, this means that when a person closes the page and then opens it a few seconds later, two sessions would be counted, while someone who leaves the page open all day would only have one session counted. Since we are interested in the continuous periods of active engagement with Facebook, we need a different definition of a session. Hence, we define a session as a series of consecutive interactions without a break longer than 10 minutes. In other words, a session consists of all interactions that are within 10 minutes of the previous interaction, as illustrated in Figure 1. Earlier research has shown that 10-15 minute is an appropriate threshold for constructing sessions [10]. Moreover, using different values left the substantive results of this paper unchanged.

4. BEHAVIORAL CHANGES DURING THE SESSION

We demonstrate that people change their content consumption behavior over the course of a session. We explore how different factors, including age and content type, affect these behavioral changes. In the first subsection, we focus on all types of content that appear in the News Feed except videos, which we consider separately because the time spent on videos is measured differently.

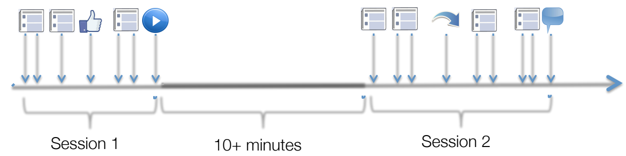

Figure 1: An illustration of a session. The timeline contains a series of interactions a person has on Facebook, which includes reading posts, liking them, playing videos, and commenting on posts. A period between consecutive interactions lasting longer than 10 minutes represents a break between sessions.

4.1 Reading Stories

How much time do people spend reading stories in their News Feed, and how does this behavior change over the course of a session? To answer these questions, we calculate the average time the users view stories as a function of time since the beginning of the session. A potential source of bias is our definition of a session, which gives rise to a data censoring problem, where a person who starts reading a story one minute before the end of the session will by definition spend at most one minute on the story, while a person who starts reading the story at the very beginning of the session could spend up to 10 minutes reading it. As a result, we would observe the average time spent consuming content decrease towards the end of the session. For a fair comparison, we do not consider the stories that take longer than one minute to read in the average of time spent on stories. These stories are a small portion (7%) of the entire data set.

Another pitfall for analysis of behavioral data is that people are heterogeneous: those who have longer sessions may be different from people who have shorter sessions. For example, (as we show below) people who have shorter sessions go faster through the stories in their feed than people engaged in longer sessions. Averaging behaviors over such heterogeneous populations could produce spurious correlations. To control for heterogeneity, we separate sessions by their length and analyze the behavior of a more homogeneous population of people who have sessions of a specific length (e.g., sessions that are 10 minutes long).

Figure 2(a) shows how the average time people spend reading stories varies as a function of time since the beginning of the session (time spent is normalized by the maximum time spent across all sessions). These data are for people accessing their News Feed through a web browser. The plot includes sessions of length 10, 20, 30, and 40 minutes, which all show a similar trend: *people read stories in their feed faster as the session progresses*. As can be seen in Figures 2(b) and 2(c), people who read News Feed stories on mobile devices, such as iOS and Android devices, have a very similar pattern of behavior.

It is unlikely that behavioral changes occurring over the length of a session are the result of differences in content relevance. Because we restrict our analysis to people who view stories in a (reverse) chronological order of the time they were posted on Facebook, it is unlikely that the length or interestingness of stories is correlated with their position in the feed. In the context of reading the News Feed, the findings suggest that as people consume content, they devote less and less time to each item. One explanation for the decrease in time spent on stories over the course of a session is that as people get closer to the end of the session, they are more likely

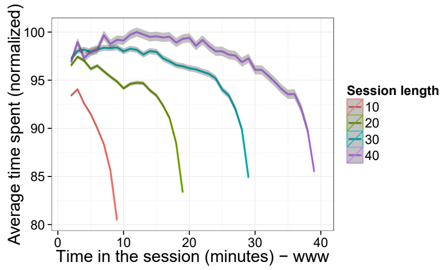

(a) Web

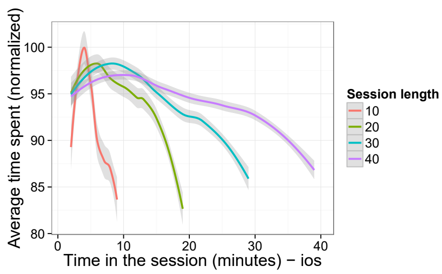

(b) iOS

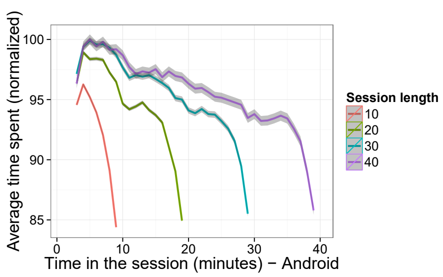

(c) Android

Figure 2: Change in time spent per story given the time in the session along with the 95% confidence interval.

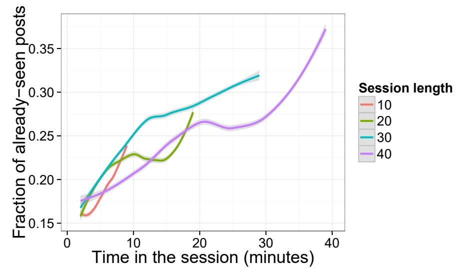

Figure 3: Change in fraction of stories that have been viewed earlier and are not new for the web users.

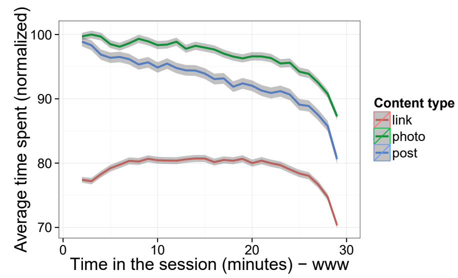

Figure 4: Change in time spent on different content types given the time in the session (web users) along with the 95% confidence interval.

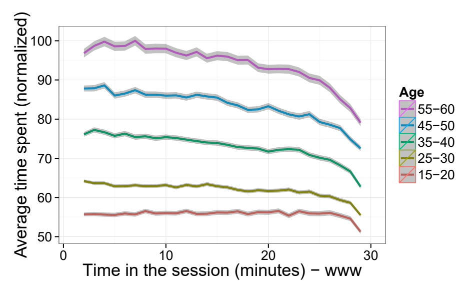

Figure 5: Change in time spent on stories for people with different age given the time in the session (web users) along with the 95% confidence interval.

to see a story they have seen before; hence, they spend less time on it (Figure 3).

In addition to the drop in time per story over the course of a session, we also observe that *people read stories faster during shorter sessions, already starting from the beginning of that session*. As we will show in this paper, the activities taking place during the first minute of a session are some of the more predictive features of session length.

The figures also show a precipitous drop at the very end of the session. We speculate that this pattern is common to people consuming the News Feed in the "most recent" configuration, where they reach the point where they encounter content they have previously consumed, rapidly scroll through a few more stories to ensure they have really reached the end of new content, and then end their session.

Next we examine the impact of different factors, such as content type, on session-level behavioral changes. We only present results for 30 minute sessions on web browsers, but the trends for other session lengths and interfaces for reading the News Feed are qualitatively similar.

Content type.

We start by considering the type of stories people consume, differentiating between photos, links to external content, and textual posts. Intuitively, different types of content require different amounts of mental effort: e.g., most people find it easier to look at a photo than read a textual post. This may cause the consumption of some types of content to be less affected by behavioral changes than others. As Figure 4 shows, in the first minute of a session, people spend almost the same amount of time viewing photos as reading textual posts, but in the last minute of the session they spent 9% more time on photos compared to textual posts. Links to external content show a smaller drop over the course of the session.

Age.

Next, we examine how age relates to session-level behavioral changes. As Figure 5 shows, age has a striking effect on the average time spent reading each story. First, older people read stories more slowly. The relative difference is as high as 80% between 15-20 year olds and 55-60 year olds. Second, and more interestingly, the behavior of older people changes more over the course of a session than for younger people, with the time spent per story experiencing a sharper drop. For the youngest age group, the average time spent

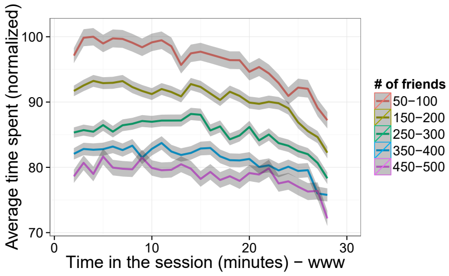

Figure 6: Change in time spent on stories for people with different number of friends given the time in the session (web users) along with the 95% confidence interval.

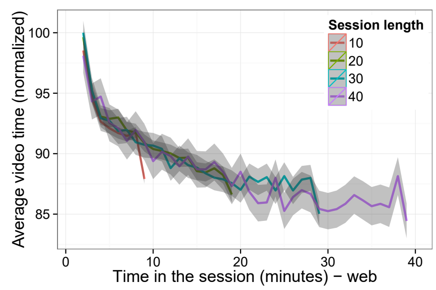

Figure 7: Change in time spent viewing videos given the time in the session (web users).

reading stories remains nearly constant over a session, though much shorter than for the older age groups. A similar trend has been found in email behavior, where older people take much longer to reply to an email compared to teenagers [13].

Number of friends.

Figure 6 compares content consumption patterns of people with many friends to those with fewer friends. People with fewer friends spend more time reading each story in general, compared to people with many friends. The slower rate of content consumption may be due to the fact that they have fewer new stories in the News Feed, so they do not need to rush through their feed to read all the stories in their limited time. Alternatively, the people with fewer friends may belong to a different population that is less familiar with the interface or generally consumes content at a slower rate. Our second observation is that people with fewer friends experience a bigger behavioral change over the course of a session compared to people with more friends: those with few (50–100) friends experience a 14% speed up in their content consumption rate between the beginning and end of a (30 minute) session, while those with many (450–500) friends experience a 9% speed up. Highly connected people interact with the larger volume of content they receive from their friends by spending less time on each story. They also do not change their behavior as much as people with fewer friends.

Time of day.

Finally, we consider the effect of the time of day on content consumption. Earlier research has shown that people's behavior changes over the course of the day, and that this is best explained by people having higher levels of energy in the morning. For example, the "morning morality effect" exists because people have higher moral awareness and self-control in the morning [14]. In the online world, people reply to a higher fraction of emails and reply to emails faster in the morning than in the evening [13]. To test the time of day effect, we consider sessions that started at different times of the day (8 am, 12 pm, 4 pm, and 10pm). People spend relatively more time to read posts in the morning (8 am) and late at night (10 pm) compared to noon (12 pm) and late afternoon (4 pm), which might be explained by the fact that most of the people are at home in early morning and late and night. Overall, there is little difference in the rate at which behavior changes over the session.

In summary, we demonstrated changes in behavior over the course of a session, with people spending less and less time on each story as they go through the feed faster and faster. This effect is more significant for some types of content: for example, textual posts, which presumably require a greater effort, show a larger decline than other content, such as photos. Age plays a considerable role in the observed effect: spending more time continuously in a session has a stronger effect on the time spent per story in older populations compared to younger ones.

4.2 Viewing Videos

Next, we analyze video viewing behavior during a session. Since video viewing time is measured as the duration of time the video plays, it gives a similar but more refined view on how people allocate their time to different types of content over the course of a session.

We observe changes in the time spent on viewing videos during a session. Following the analysis described in the previous section, we group together sessions of the same length and calculate the average amount of time spent watching videos at any minute during the session. Figure 7 shows people spend less and less time watching videos over the course of a session. However, the drop is about 5% smaller than the drop for reading stories (Figures 2 and 4). Therefore, video viewing behavior changes less during a session compared to other kinds of content, and as a result, people tend to watch relatively more videos later in the sessions.

5. PREDICTING BEHAVIOR

The patterns described in the previous section can be used to better predict people's behavior, including their session length, the number of stories consumed, and return time based on current and past behaviors.

In this section, we use the whole month of data, instead of a single day, because we want to capture people's typical behavior and leverage that for making the predictions. We use the first three weeks of the data for training, and the last week for testing. In this way, we do not use any future information in the predictions.

5.1 Session Length

We show that user activity during just the first minute of a session helps predict the ultimate session length. We frame the prediction task as a classification problem with three classes: short sessions between (1,5] minutes, medium sessions between (5,15] minutes, and long sessions, which are longer than 15 minutes. Short sessions include about 36% of all sessions, medium sessions include 37%, and long sessions account for the remaining 26% of the sessions. Note that using only the first minute of the session to predict session length is a hard task compared to other framings of the problem, such as predicting whether the session time will double.
We include a variety of features in the prediction task:

- **User characteristics:** age, gender, location, number of friends, number of days on Facebook, language, number of days active on the last 7 and 30 days, number of subscribers and subscriptions (10 features). Subscribers are users who follow updates from another user (usually celebrities), and it does not need the approval of the other user, so it is a directed relationship.

- **Session activity: features from the activity of the user in the first minute**
  - Mean, median, maximum, standard deviation, 10th, 25th, 75th, and 90th percentiles of time spent on reading posts, watching videos, and creating a post in the first minute of a session (24 features).
  - Number of different interactions during the first minute of a session: number of likes, comments, stories viewed, video playbacks, posts removed, shares, clicks, and reshares (8 features).
  - Time of the day (1 feature)
  - Time since the previous session (1 feature)
  - Number of notifications at the beginning of the session (1 feature)
- **User history: features from activity of the user in the first three weeks of the data**
  - Mean, median, maximum, standard deviation, 10th, 25th, 75th, and 90th percentiles of length of the session in minutes, number of stories viewed, and return times in the training data (24 features).
  - Counts of short, medium, and long session lengths; small, medium, and large number of stories viewed; and short, medium, and long return times in the training data (9 features).

We use the C5.0 classifier, which is a decision tree based classifier with feature selection [18]. We used a temporal split of data for training (75%) and testing (25%). For predicting the length of the session, we achieve F1 score of 0.44 and accuracy of 48.3%, which is 30.5% relative improvement over the majority vote baseline of 37.0%. Majority vote baselines always predict the majority class for each user based on the training data. We also consider a baseline that uses the empirical distribution of session length classes from the training data for each user, which then probabilistically picks one of the classes from that distribution for testing. This baseline achieves 29.3% accuracy, which is worse than the majority vote baseline; this might be because median is a typically more robust indicator of behavior than the mean user behavior.

To find the most predictive features, we rank features based on their information gain. The top three features are mean time spent reading stories, mean session length, and mean time spent on creating a post. It is perhaps unsurprising that the user history features are among the most informative features, as past behavior is often the best predictor of future behavior. However, the most informative feature (and two of the top five most informative features) come from the person's behavior in the first minute in the session. The most predictive feature is the mean time spent reading stories. This suggests that scrolling quickly through stories at the beginning of a session is a very strong signal that the person will have a short session (also seen in Figure 2).

To better understand the predictive power of different features, we divide our features into three groups and use the features in each group alone to predict session length. The first group includes features that are characteristics of the users, such as age and gender. These features will have the same value for different sessions

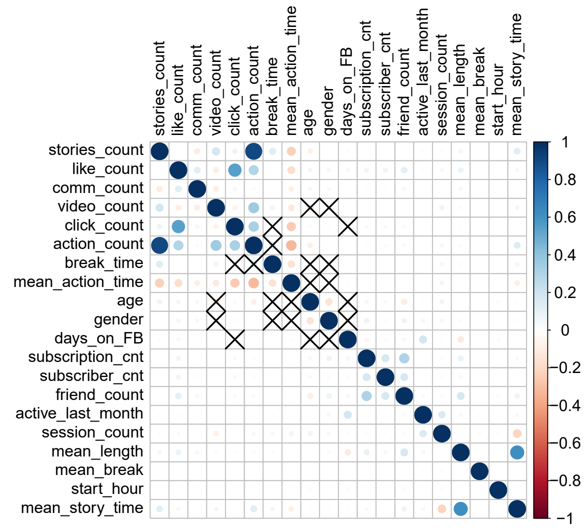

Figure 8: Correlation coefficient between the features. The color of the dots represent the value of the correlation and the size represents the absolute value. Boxes with the cross do not have a statically significant correlation.

Table 1: Prediction accuracy using different sets of features.

Prediction	Feature Group	Accuracy
Session Length	User characteristics	37.4%
	Session activity	42.5%
	User history	43.1%
Number Stories Viewed	User characteristics	37.6%
	Session activity	39.4%
	User history	46.6%
Return Time	User characteristics	36.5%
	Session activity	82.8%
	User history	42.8%

of the same user. The next group includes all features that are extracted from user activity in that particular session, e.g., mean time spent on stories. The third group includes features that are related to the earlier behavior of the user, e.g., the mean session length. Table 1 shows the accuracy of the prediction using the features in each of these groups. The user characteristics features are the least predictive, while the person's earlier Facebook browsing and their behavior at the start of the session are more predictive.

To calculate the importance of each feature individually, we first remove the correlated features (Figure 8) and then run a logistic regression over the data after normalization. Table 2 shows the ranking of the variables that have statistically significant coefficients. We find that the number of actions taken in the first minute is the top feature in ranking according to logistic regression and suggests that the number of actions users take in the first minute reflects over the period of the session. Next features are number of friends, and number of sessions that the user had in the training period. Due to privacy concerns, we are not able to share the coefficients for each feature.

5.2 Number of Stories Viewed

Next, we predict how many stories the person will view in the current session. We use the same approach and features as in the previous prediction, using the first minute of activity in the ses-

Table 2: Result of logistic regression on the independent variables for the length of the session. *** $p-value < 0.001$, ** $p-value < 0.01$, * $p-value < 0.05$

$Rank$	$Variable$	$p-val$
1	Total # of interactions	***
2	# of friends	*
3	# of sessions	***
4	Mean length of sessions	***
5	Previous break length	***
6	Mean interaction time	***
7	Time of day session started	***

Table 3: Result of logistic regression on the independent variables for the number of stories read in the session. *** $p-value < 0.001$

$Rank$	$Variable$	$p-val$
1	Mean interaction time	***
2	# of sessions	***
3	# days on Facebook	***
4	Mean session length	***
5	Age	***
6	Previous break length	***
7	Total # of interactions	***

sion. In this case we create three classes for the number of stories viewed in the session using thresholds that result in roughly balanced classes (classes include 35.1%, 32.6%, and 32.3% of data). For predicting the number of stories viewed in a session, our classifier achieves an F1 score of 0.48 and accuracy of 49.7%. This is a 41.6% improvement over the majority vote baseline, with an accuracy of 35.1%, and a 42.2% improvement over the prediction using the same probabilities that the user had in the training data, which has an accuracy of 35.3%. Our classifier is able to predict how many stories the person will consume in a session, which could be used to determine the amount of content to cache prior to the user's session.

Using only the user history features gives us the highest accuracy, compared to using features from the other two groups individually (Table 1).

Similar to the session length prediction problem, we run a logistic regression on the independent variables to find the role of each feature in the prediction. Table 3 shows the result of the regression: mean time spent on different interactions, number of sessions, and the number of days that the user has been on Facebook are the top three features for the prediction of number of stories read in the session.

5.3 Return Time

In addition to predicting the length of the session and number of stories consumed, we can also predict when the person will return to Facebook. Similar to the previous predictions, we use two thresholds that result in roughly equal-sized classes (each class includes 33.3% of the data). We use the same features, and at the end of each session we predict the time to the the next session. Using top features, our classifier achieves F1 score of 0.79 and very high accuracy of 79.0%, which is significantly higher than the other predictions and is a 137.2% relative improvement over the majority vote baseline. Interestingly, this accuracy is achieved by using only four features from the feature selection algorithm: the length of the session that just ended, mean time spent on interactions, median return time, and number of posts a person has modified (such as modifying the caption on a photo). Also, using only the length of the session that just ended achieves 73.5% accuracy, which is considerably higher than the baseline. If we use the more complex baseline, considering the history of the user, we achieve a much lower accuracy of only 36.7%.

Table 4: Result of logistic regression on the independent variables for the break length. *** $p-value < 0.001$, ** $p-value < 0.01$, * $p-value < 0.05$

$Rank$	$Variable$	$p-val$
1	Gender	*
2	Age	***
3	Mean session length	***
4	# of sessions	***
5	Time of day session started	***
6	Total # of interactions	***
7	# active days in last 30 days	***
8	Mean interaction time	***

Top three most predictive features are session length, mean time spent on interactions, and number of interactions. It's surprising that the length of the session is such a strong predictor of the length of the break, since the the length of the break is not predictive of the length of the session. In other words, people who stay longer on Facebook tend to take longer to return, but people who return to Facebook after a long break do not necessarily stay on Facebook longer. Finally, if we group the features, we observe that session activity has more predictive power compared to features extracted from user's earlier behavior (Table 1). Since our classifier uses historical information from users' earlier behavior, we cannot predict new users' behavior with the same accuracy. This problem, i.e. cold start problem, is common in recommender systems. One way to mitigate this problem is to use information from users with the same characteristics to replace missing features.

Predicting return time with such high accuracy can be extremely useful in caching the content to a users' mobile phone, by having the data ready for browsing before the user starts using the application. This could greatly improve user experience, especially in areas with poor network connectivity.

We also run a logistic regression to find the role of each feature. Gender, age, and the mean session length are the top three features in the prediction of break length (Table 4).

Table 5 summarizes all the prediction results.

6. SESSION TYPE

We have shown that regularities in user activity on Facebook can be used to predict the length of a session, how much content people will consume, and how long a break they will take. There are, however, interesting additional differences between sessions. Specifically, we show that people have different session types: sessions in which they "like" many posts (like sessions), sessions in which they tend to leave many comments (comment sessions), and sessions in which they prefer to read stories (consumption sessions).

For the analysis of sessions types, again we consider the whole dataset. For each user, we consider the distribution of likes and comments given and number of stories read across different sessions. To quantify how user activity is distributed across different sessions, we calculate the Gini coefficient of each distribution for the user. Gini coefficient measures the inequality of a distribution. Gini coefficient of zero means perfect equality: all sessions have the same number of likes or comments. On the other hand, a Gini

Table 5: Summary of the prediction results. Accuracy: percentage of correctly classified samples. Majority vote: always predicting the largest group, or predicting randomly (same group sizes).

Prediction	Majority vote baseline (random classifier)	Probabilistic baseline	Our classifier	Absolute improvement	Relative improvement	F1
Session length	37.0%	29.3%	48.3%	11.3%	30.5%	0.44
Number of stories	35.1%	35.2%	49.7%	14.6%	41.6%	0.48
Return time	33.3%	36.7%	79.0%	45.7%	137.2%	0.79

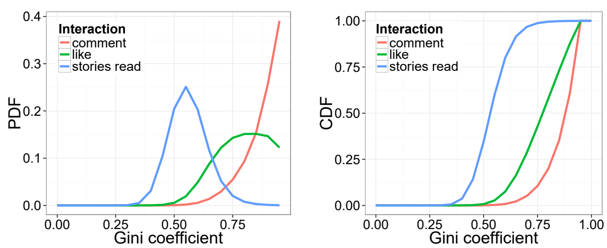

(a) PDF (b) CDF

Figure 9: Distribution of Gini coefficient of number of stories read, likes, and comments.

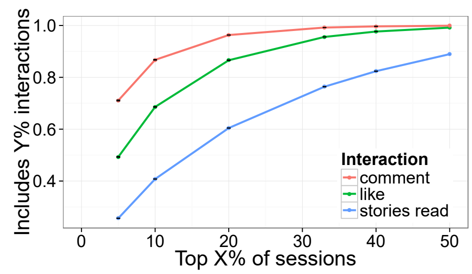

Figure 10: Small fraction of sessions include much larger portion of likes and comments, compared to stories read. 95% confidence intervals are shown, but too small to be seen, due to extremely small variance.

coefficient of one means perfect inequality: one session includes all the likes or comments, while the rest of the sessions have none.

Figure 9 shows the PDF and CDF of the distribution of Gini coefficients for the users. If the users had proportional number of likes and comments to the number of stories read in each session, the distribution of Gini coefficient will overlap with that of the number of stories read. However, the figure shows that likes and comments generally have higher Gini coefficients, and are, therefore, much more unevenly distributed compared to the stories read by the users. This means that users tend to like and comment more in some sessions, rather than distribute them evenly across all sessions and this disproportionate distribution of likes and comments cannot be explained by users having short and long sessions.

Even though we are considering a month of data, some people might have a very small number of likes and comments, and any distribution of these across the sessions would result in high inequality (i.e., high Gini coefficient). To see if small numbers are causing the high inequality for likes and comments, we randomly redistribute the likes and comments across all sessions for each user. For instance, if a user has five sessions with 7, 2, 3, 0, 1 likes in them, then the 13 total likes are randomly distributed in five bins, which might result in a distribution like 4, 2, 3, 1, 3 and the Gini coefficient is calculated for the new distribution. The randomly distributed likes have a smaller Gini coefficient (median: 0.41) compared to the non-randomized stories read (median: 0.56), while the randomized comments still have a slightly higher Gini coefficient (0.60). This suggests that perhaps some of the inequality in comments is because commenting is a relatively infrequent activity, though clearly the unshuffled comments are much more skewed than the shuffled ones. So indeed, likes and comments are unevenly distributed among the sessions.

Another way to look at this inequality is to rank a user's sessions based on the number of stories read, liked, or commented in that session. Here we see that for sessions ranked by the number of stories read, the top 5% of sessions only include 26% of the all stories read by the users. In contrast, the top 5% most liked sessions include 49% of the likes, and the top 5% most commented sessions include 71% of the comments. Figure 10 shows the same values for different percentages; the top sessions consistently include a much higher fraction of likes and comments compared to stories read.

Interestingly, the sessions that include a high number of likes are different from the sessions that have many comments. If we consider the top 5% of sessions with the most number of likes and top 5% with the most number of comments for each user, only 27% of like sessions are also comment sessions, on average (and vice versa, since we have the same number of like and comment sessions). The low overlap between like and comment sessions show that people predominantly perform certain types of actions in a session, and we are not just detecting sessions in which the person is highly engaged with Facebook.

In short, users have different session types that have many more likes and comments compared to their other sessions. In future research, it would be interesting to characterize these sessions further based on variables such as time of day, interface, and demographics, to see how easily they can be predicted.

7. CONCLUSION

We analyzed a large data set of user activities on Facebook, comprised of the interactions people have with the content their friends shared. These interactions can be divided into sessions, or periods of activity without a break longer than 10 minutes. Once segmented into sessions, content consumption shows strong regularities with predictable changes over the course of a session for many people. Regardless of the platform they use to consume Facebook content (web browser or mobile device), their demographic attributes, social connectivity, or the time to day they are active, people manifest similar behavioral changes: as the session progresses, they spend relatively less time viewing a story or video, and preferentially shift their attention away from reading stories and more towards viewing photos and videos. There were also strong differences between short and long sessions. People spent less time consuming content during shorter sessions, a pattern that was already evident at the start of the session. While our work does not address the origins of these behavioral changes, the fact that we see them in almost all user populations suggests a fruitful area of future research that delves into factors affecting differences in people's content consumption and interaction between and within sessions.

We leveraged observed behavioral regularities to predict the length of a session, how much content people will consume over the course of a session, and when they will return to Facebook. While a person's past Facebook usage offers good indicators for predicting future behavior, surprisingly, the first minute of activity was also a very good predictor of session behavior. In fact, it was the most predictive feature in our models, followed by historical features which included the average session length from the previous day and the average time spent on all interactions. These kinds of predictions could potentially be used to improve user experience, e.g. by caching content based on when and how much an individual is likely to consume it, especially in areas where internet connectivity may be poor.

Acknowledgements

The work was funded in part by the Army Research Office under contract W911NF-15-1-0142.

8. REFERENCES

[1] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user behavior in online social networks. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 49–62. ACM, 2009.

[2] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user navigation and interactions in online social networks. *Information Sciences*, 195:1–24, 2012.

[3] M. Daoud, L. Tamine-Lechani, M. Boughanem, and B. Chebaro. A session based personalized search using an ontological user profile. In *Proceedings of the 2009 ACM symposium on Applied Computing*, pages 1732–1736. ACM, 2009.

[4] C. Eickhoff, J. Teevan, R. White, and S. Dumais. Lessons from the journey: a query log analysis of within-session learning. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 223–232. ACM, 2014.

[5] S. A. Golder and M. W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, 2011.

[6] S. A. Golder, D. M. Wilkinson, and B. A. Huberman. Rhythms of social interaction: Messaging within a massive online network. In *Communities and technologies 2007*, pages 41–66. Springer, 2007.

[7] K. Goševa-Popstojanova, A. D. Singh, S. Mazimdar, and F. Li. Empirical characterization of session–based workload and reliability for web servers. *Empirical Software Engineering*, 11(1):71–117, 2006.

[8] N. Grinberg, P. A. Dow, L. A. Adamic, and M. Naaman. Changes in engagement before and after posting to facebook. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 564–574. ACM, 2016.

[9] N. Grinberg, M. Naaman, B. Shaw, and G. Lotan. Extracting diurnal patterns of real world activity from social media. In *ICWSM*, 2013.

[10] D. He and A. Göker. Detecting session boundaries from web user logs. In *Proceedings of the BCS-IRSG 22nd annual colloquium on information retrieval research*, pages 57–66, 2000.

[11] R. Jones and K. L. Klinkner. Beyond the session timeout: automatic hierarchical segmentation of search topics in query logs. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 699–708. ACM, 2008.

[12] K. Kapoor, K. Subbian, J. Srivastava, and P. Schrater. Just in time recommendations: Modeling the dynamics of boredom in activity streams. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 233–242. ACM, 2015.

[13] F. Kooti, L. M. Aiello, M. Grbovic, K. Lerman, and A. Mantrach. Evolution of Conversations in the Age of Email Overload. In *Proceedings of the 24st International World Wide Web Conference (WWW'15)*, Florence, Italy, May 2015.

[14] M. Kouchaki and I. H. Smith. The morning morality effect the influence of time of day on unethical behavior. *Psychological science*, page 0956797613498099, 2013.

[15] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. ACM, 2009.

[16] R. D. Malmgren, D. B. Stouffer, A. E. Motter, and L. A. Amaral. A poissonian explanation for heavy tails in e-mail communication. *Proceedings of the National Academy of Sciences*, 105(47):18153–18158, 2008.

[17] J. G. Oliveira and A.-L. Barabási. Human dynamics: Darwin and einstein correspondence patterns. *Nature*, 437(7063):1251–1251, 2005.

[18] R. Quinlan. Data mining tools see5 and c5. 0. 2004. https://www.rulequest.com/see5-info.html.

[19] D. E. Rose and D. Levinson. Understanding user goals in web search. In *Proceedings of the 13th international conference on World Wide Web*, pages 13–19. ACM, 2004.

[20] O. Shoukry, M. Abd El-Mohsen, J. Tadrous, H. El Gamal, T. ElBatt, N. Wanas, Y. Elnakieb, and M. Khairy. Proactive scheduling for content pre-fetching in mobile networks. In *Communications (ICC), 2014 IEEE International Conference on*, pages 2848–2854. IEEE, 2014.

[21] O. K. Shoukry and M. B. Fayek. Evolutionary content pre-fetching in mobile networks. In *2013 12th International Conference on Machine Learning and Applications*, volume 1, pages 386–391, Dec 2013.

[22] B. R. Smith, G. D. Linden, and N. K. Zada. Content personalization based on actions performed during a current browsing session, Feb. 8 2005. US Patent 6,853,982.

[23] H. Song, C. Min, J. Kim, and Y. I. Eom. Usage pattern-based prefetching: quick application launch on mobile devices. In *Computational Science and Its Applications–ICCSA 2012*, pages 227–237. Springer, 2012.

[24] M. Spiliopoulou, B. Mobasher, B. Berendt, and M. Nakagawa. A framework for the evaluation of session reconstruction heuristics in web-usage analysis. *Informs journal on computing*, 15(2):171–190, 2003.

[25] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Commun. ACM*, 53(8):80–88, Aug. 2010.

[26] Y. Wang, X. Liu, D. Chu, and Y. Liu. Earlybird: Mobile prefetching of social network feeds via content preference mining and usage pattern analysis. In *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 67–76. ACM, 2015.

[27] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 177–186. ACM, 2011.