Twitter Session Analytics: Profiling Users' Short-Term Behavioral Changes

Farshad Kooti[1(✉)], Esteban Moro[2], and Kristina Lerman[1]

[1] USC Information Sciences Institute, Marina Del Rey, USA
kooti@usc.edu
[2] Universidad Carlos III de Madrid, Madrid, Spain

Abstract. Human behavior shows strong daily, weekly, and monthly patterns. In this work, we demonstrate online behavioral changes that occur on a much smaller time scale: minutes, rather than days or weeks. Specifically, we study how people distribute their effort over different tasks during periods of activity on the Twitter social platform. We demonstrate that later in a session on Twitter, people prefer to perform simpler tasks, such as replying and retweeting others' posts, rather than composing original messages, and they also tend to post shorter messages. We measure the strength of this effect empirically and statistically using mixed-effects models, and find that the first post of a session is up to 25 % more likely to be a composed message, and 10–20 % less likely to be a reply or retweet. Qualitatively, our results hold for different populations of Twitter users segmented by how active and well-connected they are. Although our work does not resolve the mechanisms responsible for these behavioral changes, our results offer insights for improving user experience and engagement on online social platforms.

1 Introduction

Understanding people's online behavior can motivate the design of human-computer interfaces that enhance user experience, increase engagement, and reduce cognitive load. Until recently, most of the research in this area focused on web search and browsing. Researchers found it useful to segment online activity into sessions, defined as periods of time that the user is actively engaged with the platform and usually has a single intent [15,17]. For example, Kumar and Tomkins found that about half of all web page views during a typical session are of inline content, one-third are communications, and the remaining one-sixth are search [20]. Search sessions have also been studied on Twitter to compare search on Twitter and web. Researchers found that search sessions on Twitter tend to be shorter and include fewer queries compared to web search [28]. Similarly, Benevenuto et al. analyzed activity sessions on an online social network aggregator to understand how frequently and for how long people use different social networking platforms, and what sequence of actions they take during a session [16].

In this paper, we carry out a study of user activity sessions on Twitter to document short-term behavioral changes occurring over the course of a single session. Similar to earlier studies of web search, we segment the time series of an individual's activity on Twitter into sessions, where each session is a series of consecutive interactions—tweeting, retweeting, or replying—without a break longer than a specified threshold. (We experimented with different ways of defining sessions and different thresholds, and our findings are qualitatively very similar with different definitions of session.) We find that most sessions are short, but there are considerable number of sessions that span hours. Despite their short duration, we find that significant behavioral changes occur over the course of a single session, with people preferring easier interactions later in the session. Specifically, people tend to compose longer tweets at the beginning of a session, and reply and retweet more later in the session, and also when there is a short time period between consecutive interactions. While Twitter population is highly heterogeneous, these patterns hold across different subsets of the population, e.g., for both highly connected and poorly connected users, as well as for highly-active and less-active users.

Earlier studies have shown strong daily, weekly, and monthly patterns in social activity. For example, Foursquare check-ins, mobile phone calls, or tweets show strong daily and weekly patterns corresponding to food consumption and nightlife [13], different social contexts [1], economical activity [21], or worldwide daily and seasonal mood variations in Twitter [9]. In this work, we find patterns that occur in far shorter time scales of only a few minutes, compared to daily and monthly patterns of earlier work. While long term patterns can be explained by the circadian cycles, work schedules, and other global macroscopic forces, the behavioral changes we study appear to be qualitatively different, arising from the individual decisions (perhaps unconscious) to allocate attention and effort. To our knowledge, this is the first demonstration of short-term behavioral changes on Twitter.

The main contributions of our work are as follows:

– We present a detailed analysis of user activity sessions on Twitter. We show that most of the sessions are very short; however, while large fraction of sessions include only one type of tweet, most of the sessions are mixture of different types of tweets (e.g., normal tweets, replies, and retweets) (Sect. 2).
– We show that later in a session people tend to perform easier or more socially rewarding interactions, such as replying or retweeting, instead of composing original tweets. Also, they tend to compose shorter tweets later in a session (Sect. 3).
– We divide people based on their characteristics, such as position in the follower graph or activity, and show that people with higher activity or more friends behave differently (Sect. 4).

Several mechanisms could explain our observations. First, deterioration of performance following a period of sustained mental effort has been documented in a variety of settings, including data entry [14] and exerting self-control [23], and led researchers to postulate cognitive fatigue [2] as the explanation. On

Twitter, as people become fatigued over the course of a session, they may switch to easier tasks that require less cognitive effort, such as retweeting instead of composing original tweets. Alternately, our observations could be explained by growing boredom or loss of motivation. It is plausible that social interactions are highly motivating, and the fact that users continue to reply to others, even when they are less likely to create original tweets, appears to indicate that they shift their effort to the more engaging tasks, such as social interactions. Still other explanations are possible, such as users' choice to strategically shift their attention to other tasks. While our work does not address the causes of these behavioral changes, our findings are significant in that they can be used to predict users' future actions, which could, in turn, be leveraged to improve user online experience on social platforms.

2   Methods

Our Twitter dataset includes more than 260 M tweets posted by 1.9 M randomly selected users and all their tweets, using Twitter's API. Twitter is known to include lots of spammers. To eliminate spammers from our dataset, we took the approach of [8] and classified users as spammers or bots based on entropy of content generated and entropy of time intervals between tweets (spammers and bots tend to have low entropy of content and tweeting time intervals).

   User online activity can be segmented into sessions, usually characterized by a single intent [15, 17]. We apply a similar idea to our Twitter data. To construct activity sessions from the time series of user's tweets, we examine the time interval between successive tweets and consider a break between sessions to be a time interval greater than some threshold. Following [17], we use a 10-min threshold. Thus, all tweets posted by a user within 10 min of his or her previous tweet are considered to be in the same session, and the first tweet posted following a time period longer than 10 min starts a new session (Fig. 1). We experimented with different time thresholds and the results remain robust. Due to the heavy-tailed distribution of inter-tweet time interval, increasing the threshold only merges a very small fraction of sessions. Figure 2 shows the probability (PDF) and cumulative (CDF) distribution of time between consecutive tweets. This distribution is very similar to the distribution of time between phone calls a person makes [25]. There is no clear cut-off and the plot drops gradually. This figure also shows that increasing the 10 min threshold to 30 min, only affects 6 % of the sessions.

   To understand sessions, we look at the distribution of session length (time interval between the first and last tweet of the session) and number of tweets posted in the session. While these distributions would change if a different time threshold was used, as explained above, the change is not significant. Most of the sessions include few tweets: 64 % of sessions include only two tweets, and only 1 % include 12 or more tweets. Moreover, sessions tend to be very short: 99 % of sessions are only 1 min long, even if we only consider sessions that include 5 tweets or more, 98 % of them are still only 1 min long.

   We also analyze the types of tweets that are posted in a session. We classify tweets into three main types:
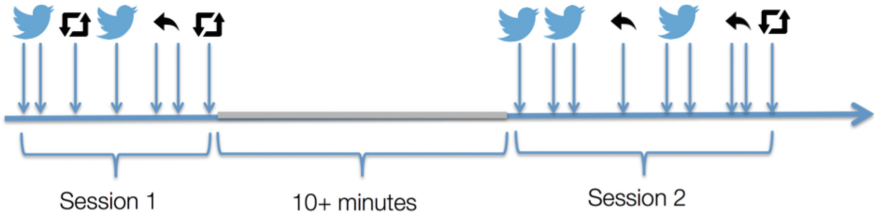
Fig. 1. Timeline of user activity on Twitter segmented into sessions. The timeline is a time series of tweets, including normal tweets, retweets, and replies. These activities fall into sessions. A period between consecutive tweets lasting longer than 10 min indicates a break between sessions.


(a) PDF   (b) CDF
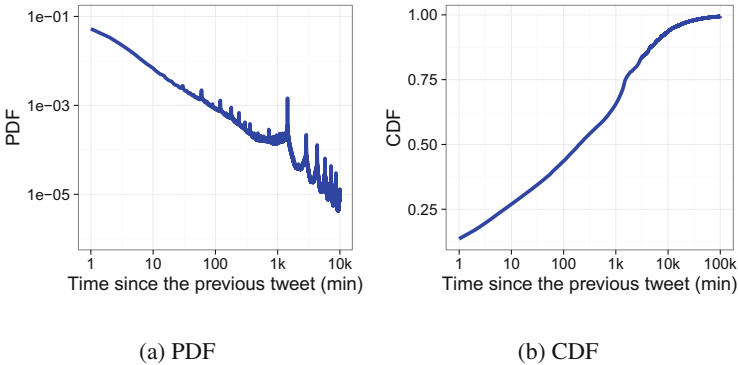
Fig. 2. Distribution of the time interval between consecutive tweets.

reply a message directed to another user, usually starting with an @mention.
retweet an existing message that is re-shared by the user, sometimes preceded by an 'RT'
normal all other tweets; typically composed tweets, which may include urls and hashtags

Considering all sessions, 59 % of sessions include only one type of tweet. This percentage is very high because a large fraction of sessions include only two tweets, so there is a very low probability of diversity. Considering only sessions that include more than five tweets, then only 35 % of the sessions include one type of tweet, 41 % include two types of tweets, and the remaining 24 % include all three types of tweets. To better understand the diversity of sessions, we consider sessions that include 10 tweets and cluster them based on the fraction of normal tweets, replies, and retweets. We use the X-means algorithm from Weka[1] that automatically detects the number of clusters. The algorithm creates three clusters, where in each cluster one type of tweet is dominant. 44 % of sessions belong to the cluster where majority of tweets are normal, 31 % are sessions with

---

[1] http://weka.sourceforge.net/doc.packages/XMeans/weka/clusterers/XMeans.html.
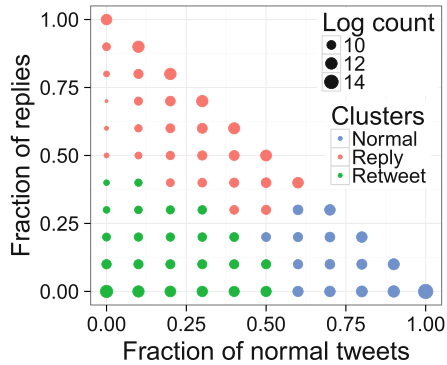
Fig. 3. Visualization of clustering of sessions using the fraction of normal tweets, replies, and retweets. (Color figure online)

many replies, and 25 % of the sessions include mostly retweets. Figure 3 shows a visualization of the sessions with each color representing a cluster and the size of dots representing the number of sessions with that fractions of tweet types. The x-axis shows the fraction of normal tweets in the session, and y-axis shows the fraction of replies in the session. Each cluster could be found in the plot by considering the fractions, e.g., the red circles belong to replies, because they have high fraction of replies, and the green circles belong to the retweet cluster, because they have low fraction of normal tweets and replies. As it is shown in the figure, these clusters are not clearly separated and there is a spectrum of sessions with different fraction of tweet types. This means there is no clear users or sessions that have a particular purpose, and most of the sessions include a mixture of different types of tweets.

3   Session-Level Behavioral Changes

In this section, we present evidence for changes in user behavior over the course of a single session on Twitter. We focus on three types of behaviors: (i) the type of the message (tweet) a user posts on Twitter, (ii) the length of the message the user composes, and (iii) the number of spelling errors the user makes. Since sessions are typically short, with the vast majority lasting only a few minutes, the demonstrated behavioral changes take place on far faster time scales than those previously reported in literature (e.g., diurnal and seasonal changes).

3.1   Time to Next Tweet

The type of a tweet a user posts depends on how much time has elapsed since the user's previous interaction on Twitter. As shown in Fig. 4, 30 % of the tweets posted 10 s after another tweet are normal tweets, whereas more than 50 % of tweets posted two minutes or more following a previous tweet are normal tweets.

In general, the longer the period of time since a user's last action on Twitter, the more likely the new tweet is to be a normal tweet. Note that we excluded tweets posted within 10 s of the previous tweet, because they are likely to have been automatically generated, e.g., by a Twitter bot. Despite the filtering, our data still contains some machine-generated activity, as evidenced by spikes at 60 s, 120 s, etc. The shorter the time delay from the previous tweet, the more likely the tweet is to be a retweet. Replies are initially similar to normal tweets: the more time elapsed since the previous tweet, the more likely the new tweet is to be a reply, but unlike normal tweets, their probability saturates and even decreases slightly with longer delays.
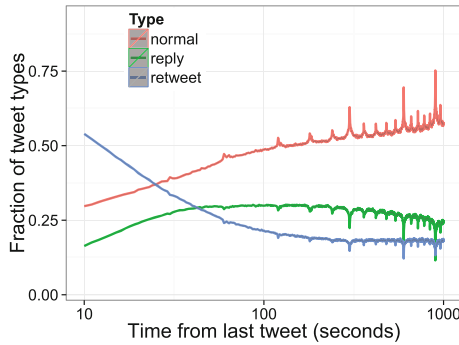
Fig. 4. Fraction of different tweet types given the time from the user's last tweet.

To understand these temporal patterns, we segment a user's activity into sessions, as described in the previous section. We can characterize sessions along two dimensions: (a) the number of tweets produced during the session and (b) the length of the session in terms of seconds or minutes, i.e., the time period between the first and last tweet of the session. Each of these dimensions plays an important role in the types of the tweets that are produced during the session. For example, short sessions with many tweets are very intense, and the user may not have enough time to compose original tweets; hence, the tweets are likely to be replies. On the other hand, a long session with few tweets is more likely to include more normal tweets, because the user has had enough time to compose them. The fraction of tweets that are replies is shown in Fig. 5, which shows these trends: users are more likely to reply as sessions become longer (in time), or there are fewer tweets posted during sessions of a given duration.

We can study the behavioral change with respect to either the position of the tweet in the session or the time elapsed since the beginning of the session. Our preliminary analysis showed that the number of tweets in a session plays a more significant role compared to the time since the first tweet of the session. Hence, in the following analyses, we study changes with respect to the position of the tweet within a session and not with respect to the time since the first tweet. In general, the trends are similar but weaker if we consider the time since the first tweet of the session.
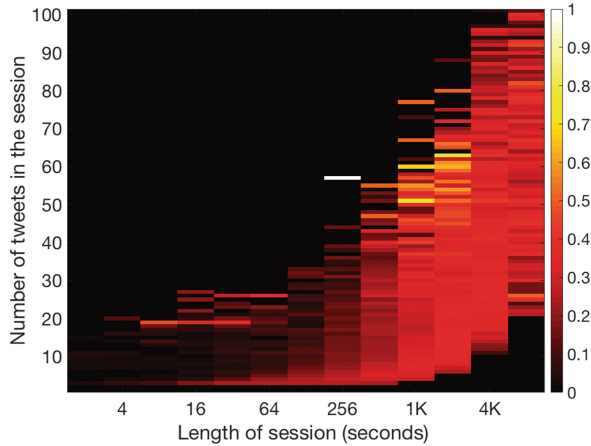
Fig. 5. Fraction of tweets that are replies posted during sessions of a given length in time and number of tweets in the session. The data was binned and only bins with more than 100 sessions are included.

3.2 Changes in Tweet Type

Next, we study the types of tweets that are posted at different times during a session. Since user behavior during longer sessions could be systematically different from their behavior during shorter sessions, we aggregate sessions by their length, which we define as the number of tweets posted. Then for each tweet position within a session, we calculate the fraction of tweets that belong to each of our three types. Figure 6 shows that tweets are more likely to be normal tweets early in a session, and later in a session, users prefer cognitively easier (i.e., retweet) or socially more rewarding (i.e., reply) interactions.

Since user population on Twitter is highly heterogenous, these observations could result from non-homogeneous mixing of different user populations. Kooti et al. show an example of this, where a specific population of users is over-represented on one side of the plot (e.g., early during a session), producing a trend that does not actually exist [18]. One way to test for this effect is through a *shuffle test*. In a shuffle test, we randomize the data and conduct analysis on the randomized (i.e., shuffled) data. If the analysis of the shuffled data yields a similar result as of the original data, then the trend is simply an artifact of the analysis and does not exist in the data. If trends disappear completely, it suggests that the original analysis is meaningful.

To shuffle the data, we reorder the tweets within each session, keeping the time interval between them the same. Figure 7 shows results of the analysis on the shuffled data. Flat lines indicate that the factions of all tweet types do not change over the course of the shuffled session. This suggests that the trends observed in the original data have a behavioral origin.

We use values in Fig. 7 as baseline to normalize the average fraction of tweets types in Fig. 6. Figure 8 shows the change in the fraction of tweet types relative
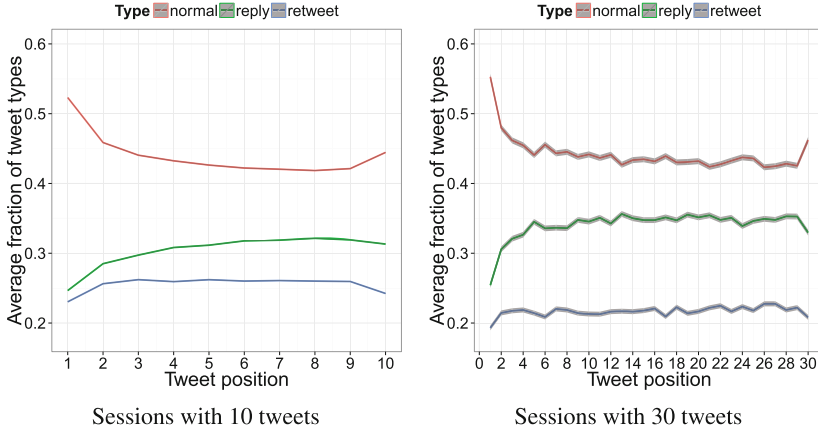
Fig. 6. Change in the fraction of tweets of each type over the course of sessions in which users posted 10 or 30 tweets.
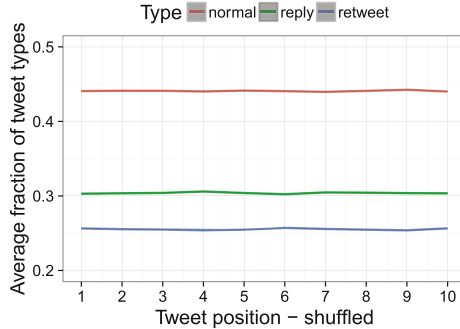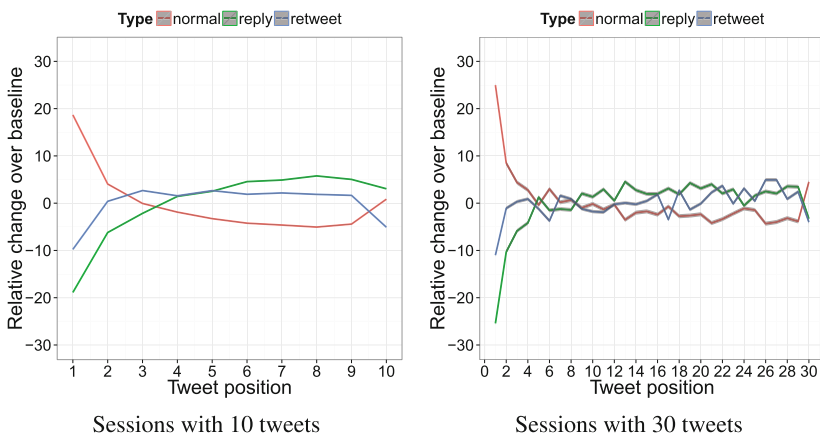

Fig. 7. Change in the fraction of tweets of each type over the course of sessions of length 10 in shuffled data.

to the baseline and clearly shows that the first tweets of a session are up to 30 % relatively more likely to be normal tweets, and 10–20% less likely to be replies or retweets. The time when a normal tweet becomes less likely than the baseline (red line crossing zero) is later during longer sessions, and it happens after ∼30 % of the tweets are posted, i.e., at the 3rd position for sessions with 10 tweets and at the 10th position in sessions with 30 tweets.

What explains the observed trends? To partially address this question, we focus on the fraction of replies. As explained above, users are more likely to reply later in a session rather than compose an original tweet. This may arise because some sessions are extended by the ongoing conversations the user has with others. To test this hypothesis, we calculate the fraction of replies at each position within the session that are in response to a tweet that was posted since the start of that session. In other words, we calculate the fraction of replies in

Fig. 8. Relative change in the fraction of tweets of each type over the course of sessions with 10 or 30 tweets.

conversations initiated during that session. Figure 9 shows this fraction: replies that are posted later in the session are much more likely to belong to an ongoing conversation. This means that some part of the trend found above could be explained by users extending their sessions to interact with others.

3.3 Change in Tweet Length

Next, we study the change in the length of tweets posted over the course of a session. We exclude retweets from this analysis, because length of the retweets does not represent the effort needed to compose them. First, we calculate the average length of the tweet at each position in the session, but there is too much variation in tweet length to produce any statistically significant trends. Instead,
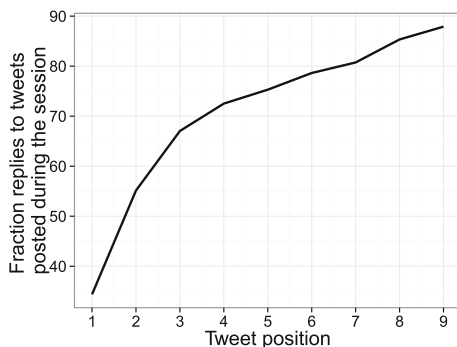
Fig. 9. Fraction of tweets that are replies to tweets posted since the beginning of the same session (for sessions with 10 tweets) .
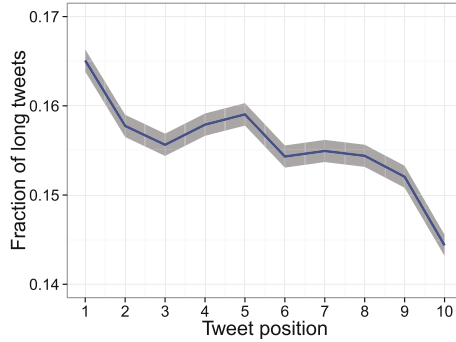
Fig. 10. Fraction of long tweets posted over the course of sessions of a given length (10 tweets). Long tweets are defined as non-reply tweets that are longer than 130 characters.

we divide tweets into long (longer than 130 characters) and short tweets (shorter than 130 characters), and measure the fraction of long tweets over the course of the session. We find a statistically significant trend, wherein tweets posted later in the session are more likely to be short, compared to tweets posted earlier in the session (Fig. 10). We choose a high threshold for the long tweets, because when a user is reaching the 140 character limit imposed by Twitter, they usually have to make an effort to shorten their tweet by rephrasing and abbreviating the message. We believe that this results in a stronger signal for analysis, compared to the situation where the user is just typing a few more characters e.g., 30 characters vs. 35 characters. To ensure that the drop in the fraction of long tweets is a real trend, we perform the shuffle test and obtain a flat line. This suggests that users are less likely to devote the effort to compose long tweets later in a session. We exclude tweets including URLs and repeat the analysis again, and we achieve very similar results. Similarly, considering only normal tweets and replies results in the similar trend.

3.4 Change in the Number of Spelling Mistakes

Finally, we consider the percentage of words that are spelled incorrectly in a tweet. Earlier studies have shown that when people are tired their judgment is impaired [3], and it is harder for them to solve problems correctly [14]. We hypothesize that we can observe this effect in terms of number of spelling errors that users make. To this end, for each tweet we calculate the percentage of words that are spelled incorrectly (i.e., typos) and calculate the average percentage of typos at each tweet position in a session. We exclude retweets, non-English tweets, and punctuations and use a dictionary that includes all forms of a word, e.g., including the past tense of the verbs and the plural of the nouns.

Figure 11 shows that there is a small but statistically significant increase in the percentage of typos made in tweets over the course of a session. This percentage rises quickly initially, but saturates later in the session. Overall, there
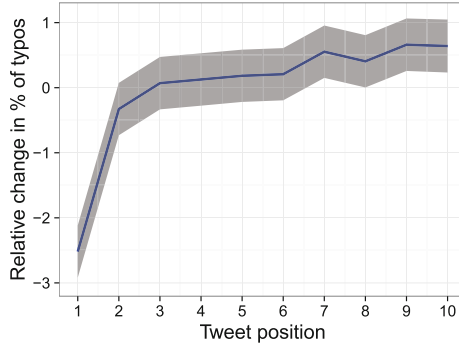
Fig. 11. Percentage of change of spelling errors made in tweets over the course of session relative to shuffled data.

is a 3% relative increase in the probability of making a spelling mistake later in the session, compared to first tweets of the session. The same trend exists for replies and normal tweets when considered individually.

3.5   Modeling

The results presented above strongly suggest that tweeting behavior changes over the course of a session. To make these findings more quantitative, we model the trends statistically. One challenge for statistical analysis is that the data samples are not independent, as we have multiple sessions from the same user. In addition, there is significant heterogeneity among the users, with some users posting mostly normal tweets, while the others mostly retweeting. As a result, our conclusions, which are based on data aggregated over the entire population, could be affected by the heterogeneous mixture of different populations (Simpson's paradox). To resolve this issue, we model the tweeting activity using mixed-effects models, which consider the individual differences.

The mixed-effects models include two main components: (i) fixed effects, which are constant across different user populations, e.g., the index or position of the tweet in the session, and (ii) random effects, which vary across different users, e.g., reflecting user's preference to post tweets of a particular type. The random effect enables us to consider individual differences among users to identify the role of the fixed effects.

We model each tweet type independently as a binary response. The model determines if a tweet is a particular tweet type given the position of the tweet in the session, the session length, and considering the user who has posted the tweet. This model can be written as *tweet type* $\sim 1 + $ *tweet index* $+$ *session length* $+$ $(1|user)$. We represent the intercept of the model by 1, and the next two terms are the fixed effects that we are interested in, and finally the particular user is also considered. In modeling the normal tweets, the coefficient of the tweet index is -0.0148, meaning that tweets posted later in the session are less likely

to be a normal tweet. On the other hand, in the model for replies, the tweet index coefficient is $+0.0149$, confirming our earlier findings and showing that tweets that are posted later in the session are more likely to be a reply. For retweets, the index coefficient is -0.0001, which is very small and negative, meaning retweeting is slightly less likely later in the session. This is due to the strong over-representation of replies later in the sessions, and if we consider only normal tweets and retweets, then the index coefficient becomes positive. The median scaled residuals for the three models are only -0.07 for modeling normal tweets, and -0.19 for modeling replies and retweets, showing that the model has a very low rate of errors.

In short, we considered the individual differences by modeling the tweet types using mixed-effects models. The results of the modeling confirmed that the results of our empirical analyses are not due to aggregating over different user population.

4   User Characteristics

In this section, we investigate how differences between users may contribute to behavioral changes. We split users based on their characteristics and carry out analysis described in the previous section within subpopulations of users.

4.1   User Connectivity

One of the main characteristics of Twitter users is the number of friends they have, i.e., the number of other Twitter users they follow. This number is highly correlated with the amount of information users receive and the number of inter-actions they have with other users. We rank users based on the number of friends and compare the session-level behavioral differences of the bottom 20 % with the top 20 %. In both cases, we measure how the fraction of tweet types change relative to the baseline, over the course of a session. Figure 12 shows that users with many friends retweet significantly more compared to users who follow few others. This is perhaps not surprising, as the well-connected users tend to receive many more tweets and have more opportunities for retweeting. These users also tend to be very active, and as users become more active, they tend to retweet more (arguably because it takes less effort). However, even though the fraction of tweet types is different in the two groups, the change over the course of a session is very similar. Therefore, we conclude that users with different numbers of friends act differently in general, but their behavior changes the same way over the course of a session. We verify that the results are not an artifact of the analysis by performing the shuffle test.

4.2   User Activity

Next, we divide users into different classes based on their activity, i.e., the rate of tweeting. We order users based on the average number of tweets in a month,
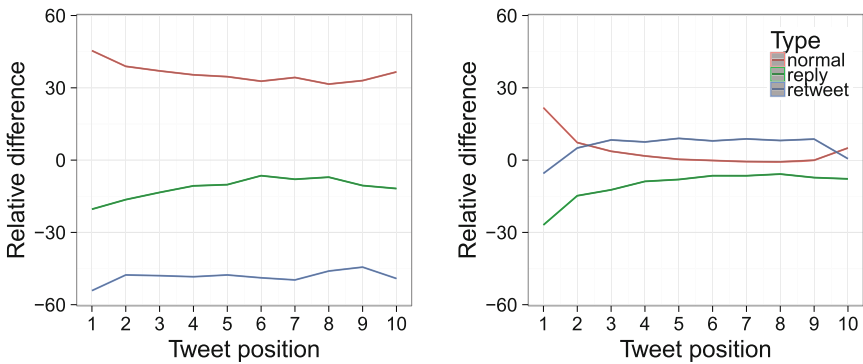
(a) Few friends ($<$ 46 friends)   (b) Many friends ($>$ 1,321 friends)

Fig. 12. Relative change in the tweet type throughout a session for users with few friends and many friends. The change is relative to shuffled sessions with 10 tweets.

and compare the top 20 % of the most active users to the bottom 20 % of the users. We find that the less active users tend to compose more original (normal) tweets, and are more likely to do it than users with most tweets. In contrast, the more active users produce many more retweets and replies, compared to users with lower levels of activity (Fig. 13). And, unlike previous analysis that divided users based on the number of friends, the change in the fraction of replies shows a higher increase for more active users. We again conduct a shuffle test to ensure that the observed effect is real.
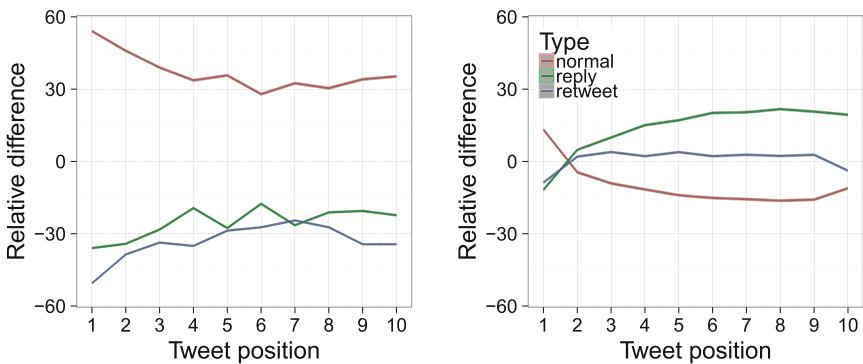
(a) Low activity ($<$ 0.93 tweets/month)   (b) High activity ($>$ 52.5 tweets/month)

Fig. 13. Relative change in tweet type throughout a session for users with low and high activity.

We conclude that part of what makes users active is their willingness to engage in social interactions on Twitter. Users extend their session to carry on conversations with others. People appear to prioritize their online activity on Twitter, and social interactions appear to be preferable, especially more active users, later in the session.

5   Related Work

Sessions of activity have shown to be an effective way to characterize people's online behavior, by segmenting a person's activity to meaningful smaller sections that are easier to study and analyze [5,24,26]. In the research community, sessions are usually constructed in two ways: a series of actions that serve a single intent [6,17], or more commonly, a period of time without a break longer than a given threshold [11,27], which is our definition of session.

Sessions have been studied extensively in context of browsing and search behavior [15,17,20]. In the recent years, sessions of activity have been also used for understanding users' behavior in online social networks. Benevenuto et al. created sessions of activity from a social network aggregator to understand users' behavior in high-level, e.g. how frequently and for how long the social networks are used [16]. On Twitter, Teevan et al. studied sessions to compare Twitter search with web search [28]. And more recently on Facebook, Grinberg et al. studied the effect of content production on length and number of sessions [12].

The changes in behavior of users over the course of a session could be attributed to fatigue or cognitive depletion. These concepts have been studied extensively in the offline world by psychologists. They have shown that there is a temporal component in cognitive performance. Mental effort makes it more difficult for people to perform cognitively demanding tasks at a later time, whether to solve problems correctly [14], make a decision [3], or exercise self-control [7,22]. The phenomenon of lower cognitive ability after sustained mental effort is generally referred to as "ego depletion" [4]. Although there have been multiple proposals for various mechanisms of ego depletion and they are still debated, there is consensus among researchers that cognitive performance declines over a period of continuous mental effort. Our study is another evidence for this phenomena.

Our study presents behavioral changes that occur on a very small time scale; only in order of minutes. Multiple studies have shown daily, weekly, monthly, and yearly patterns of activity in offline and online world: people make more donations in the mornings [19], strong daily and weekly patterns of food consumption exist in Foursquare checkins [13], there are significant seasonal patterns in communications among college students on Facebook [10], or diurnal and seasonal trends affect people's sentiment expressed on Twitter posts [9].

6   Conclusion

In this work, we analyzed user behavior during activity sessions on Twitter. We found that users engage with Twitter usually for short periods of time, what we

refer to as activity sessions, that are on the order of minutes and include only a few tweets. The tweets posted during these times tend to be diverse tweets, including original (composed) messages, retweets of others' messages, and replies to other users. Despite its short duration, users' behavior changes over the course of a session, as they appear to prioritize different types of interactions. The longer they are on Twitter, the more they prefer to perform easier or more socially engaging tasks, such as retweeting and replying, rather than harder tasks, such as composing an original tweet. This effect is quite large: at the beginning of the session, the tweets are up to 25 % more likely to be original tweets than near the end of the session.

We also found that tweets tend to get shorter later in the session, and people tend to make more spelling mistakes. All these results could be explained by people becoming cognitively fatigued, or perhaps careless due to loss of motivation. If we divide users into classes based on the number of friends they follow, or their activity level (i.e., the number of tweets they posted), we find that while these user classes behave differently in general, in terms of the types of tweets they tend to post, all classes manifest similar behavioral changes over the course of the session. While our work does not resolve the mechanisms responsible for these behavioral changes, our findings are significant in that they can be used to forecast dynamics of user behavior, which could, in turn, be leveraged to improve user online experience on social platforms.

References

1. Aledavood, T., López, E., Roberts, S.G., Reed-Tsochas, F., Moro, E., Dunbar, R.I., Saramäki, J.: Daily rhythms in mobile telephone communication. PloS One **10**(9), e0138098 (2015)
2. Baumeister, R.F., Bratslavsky, E., Muraven, M., Tice, D.M.: Ego depletion: is the active self a limited resource? J. Pers. Soc. Psychol. **74**(5), 1252 (1998)
3. Baumeister, R.F., Sparks, E.A., Stillman, T.F., Vohs, K.D.: Free will in consumer behavior: self-control, ego depletion, and choice. J. Consum. Psychol. **18**(1), 4–13 (2008)
4. Baumeister, R.F., Vohs, K.D.: Self-regulation, ego depletion, and motivation. Soc.Pers. Psychol. Compass **1**(1), 115–128. http://dx.doi.org/10.1111/j.1751-9004.2007.00001.x
5. Daoud, M., Tamine-Lechani, L., Boughanem, M., Chebaro, B.: A session based personalized search using an ontological user profile. In: Proceedings of the 2009 ACM Symposium on Applied Computing, pp. 1732–1736. ACM (2009)
6. Eickhoff, C., Teevan, J., White, R., Dumais, S.: 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, pp. 223–232. ACM (2014)
7. Gailliot, M.T., Baumeister, R.F., DeWall, C.N., Maner, J.K., Plant, E.A., Tice, D.M., Brewer, L.E., Schmeichel, B.J.: Self-control relies on glucose as a limited energy source: willpower is more than a metaphor. J. Pers. Soc. Psychol. **92**(2), 325 (2007)
8. Ghosh, R., Surachawala, T., Lerman, K.: Entropy-based classification of retweeting activity on twitter. In: Proceedings of KDD Workshop on Social Network Analysis (SNA-KDD), August 2011

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

10. Golder, S.A., Wilkinson, D.M., Huberman, B.A.: Rhythms of social interaction: Messaging within a massive online network. In: Communities and Technologies 2007, pp. 41–66. Springer (2007)

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

12. Grinberg, N., Dow, P.A., Adamic, L.A., Naaman, M.: Extracting diurnal patterns of real world activity from social media. In: CHI (2016)

13. Grinberg, N., Naaman, M., Shaw, B., Lotan, G.: Extracting diurnal patterns of real world activity from social media. In: ICWSM (2013)

14. Healy, A.F., Kole, J.A., Buck-Gengle, C.J., Bourne, L.E.: Effects of prolonged work on data entry speed and accuracy. J. Exp. Psychol. Appl. **10**(3), 188–199. http://view.ncbi.nlm.nih.gov/pubmed/15462620

15. Huang, J., Efthimiadis, E.N.: Analyzing and evaluating query reformulation strategies in web search logs. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, pp. 77–86. ACM (2009)

16. Jin, L., Chen, Y., Wang, T., Hui, P., Vasilakos, A.V.: Understanding user behavior in online social networks: a survey. IEEE Commun. Mag. **51**(9), 144–150 (2013)

17. Jones, R., Klinkner, K.L.: 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, pp. 699–708. ACM (2008)

18. Kooti, F., Lerman, K., Aiello, L.M., Grbovic, M., Djuric, N., Radosavljevic, V.: Portrait of an online shopper: understanding and predicting consumer behavior. In: Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM 2016), San Francisco, USA, February 2016

19. Kouchaki, M., Smith, I.H.: The morning morality effect the influence of time of day on unethical behavior. Psychol. Sci. **25**(1), 95–102 (2013). 0956797613498099

20. Kumar, R., Tomkins, A.: A characterization of online browsing behavior. In: Proceedings of the 19th International Conference on World Wide Web, pp. 561–570. ACM (2010)

21. Llorente, A., Garcia-Herranz, M., Cebrian, M., Moro, E.: Social media fingerprints of unemployment. PloS One **10**(5), e0128692 (2015)

22. Muraven, M., Tice, D., Baumeister, R.: Self-control as a limited resource: regulatory depletion patterns. J. Pers. Soc. Psychol. **74**(3), 774 (1998)

23. Muraven, M., Baumeister, R.F.: Self-regulation and depletion of limited resources: does self-control resemble a muscle? Psychol. Bull. **126**(2), 247 (2000)

24. Rose, D.E., Levinson, D.: Understanding user goals in web search. In: Proceedings of the 13th International Conference on World Wide Web, pp. 13–19. ACM (2004)

25. Saramäki, J., Moro, E.: From seconds to months: an overview of multi-scale dynamics of mobile telephone calls. Eur. Phys. J. B **88**(6), 1–10 (2015)

26. Smith, B.R., Linden, G.D., Zada, N.K.: Content personalization based on actions performed during a current browsing session, uS Patent 6,853,982, 8 February 2005

27. Spiliopoulou, M., Mobasher, B., Berendt, B., Nakagawa, M.: A framework for the evaluation of session reconstruction heuristics in web-usage analysis. Inf. J. Comput. **15**(2), 171–190 (2003)

28. Teevan, J., Ramage, D., Morris, M.R.: # twittersearch: a comparison of microblog search and web search. In: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, pp. 35–44. ACM (2011)