

GeoLifecycle: User Engagement of Geographical Exploration and Churn Prediction in LBSNs

YOUNG D. KWON, HKUST
 DIMITRIS CHATZOPOULOS, HKUST
 EHSAN UL HAQ, HKUST
 RAYMOND CHI-WING WONG, HKUST
 PAN HUI, HKUST and University of Helsinki

As Location-Based Social Networks (LBSNs) have become widely used by users, understanding *user engagement* and predicting *user churn* are essential to the maintainability of the services. In this work, we conduct a quantitative analysis to understand user engagement patterns exhibited both offline and online in LBSNs. We employ two large-scale datasets which consist of 1.3 million and 62 million users with 5.3 million reviews and 19 million tips in Yelp and Foursquare, respectively. We discover that users keep traveling to diverse locations where they have not reviewed before, which is in contrast to “human life” analogy in real life, an initial exploration followed by exploitation of existing preferences. Interestingly, we find users who eventually leave the community show distinct engagement patterns even with their first ten reviews in various facets, e.g., geographical, venue-specific, linguistic, and social aspects. Based on these observations, we construct predictive models to detect potential churners. We then demonstrate the effectiveness of our proposed features in the churn prediction. Our findings of geographical exploration and online interactions of users enhance our understanding of human mobility based on reviews, and provide important implications for venue recommendations and churn prediction.

CCS Concepts: • **Information systems** → *Data mining*; • **Human-centered computing** → *Ubiquitous and mobile computing*; • **Social and professional topics** → *Geographic characteristics*.

Additional Key Words and Phrases: Location Based Social Networks; User Engagement; Churn Prediction; Applied Machine Learning; Yelp; Foursquare; Swarm

ACM Reference Format:

Young D. Kwon, Dimitris Chatzopoulos, Ehsan ul Haq, Raymond Chi-Wing Wong, and Pan Hui. 2019. GeoLifecycle: User Engagement of Geographical Exploration and Churn Prediction in LBSNs. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3, Article 92 (September 2019), 29 pages. <https://doi.org/10.1145/3351250>

1 INTRODUCTION

Location-Based Social Networks (LBSNs), with the most popular being Yelp and Foursquare, have become a vital part of our society due to their assistance on users’ needs. For example, foreign tourists in San Francisco can easily find highly-reputed restaurants even if it is their first time in the city. In addition, apart from assisting ordinary users every day, LBSNs also provide essential information, in the form of datasets and APIs, for

Authors’ addresses: Young D. Kwon, HKUST, ydkwon@cse.ust.hk; Dimitris Chatzopoulos, HKUST, dcab@cse.ust.hk; Ehsan ul Haq, HKUST, euhaq@connect.ust.hk; Raymond Chi-Wing Wong, HKUST, raywong@cse.ust.hk; Pan Hui, HKUST, University of Helsinki, panhui@cse.ust.hk.

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

© 2019 Association for Computing Machinery.
 2474-9567/2019/9-ART92 \$15.00
<https://doi.org/10.1145/3351250>

researchers. Popular topics in LBSNs are: Point-of-Interest (POI) recommendation [8, 9, 88, 92], user mobility [15], privacy [14, 79, 84, 85], modeling users' behaviors [13, 19, 52, 78], and urban computing [18].

There exist two types of LBSN users, the ones that produce information (e.g., by writing a review regarding a restaurant) and the ones that consume information (e.g., by reading reviews of the restaurants in one area). The producer-type users create User-Generated Content (UGC) that can affect the quality of experience of the consumer-type users [97] and the sustainability of the services by voluntarily sharing their location-related stories in LBSNs [58]. Considering that LBSNs heavily rely on UGC and users can stop contributing at any time, it is important for LBSN platforms to attract new users and keep the existing ones [39]. Hence, understanding the *user engagement* (i.e., the desire to use an application longer and repeatedly [46]) and predicting the *user churn* (i.e., the loss of a user from a service) is essential to the maintainability of the services [23, 86, 98].

The limitations of the existing studies can be categorised into two groups: (1) While there exist many studies investigating user engagement patterns in online settings [5, 21, 31, 53, 64, 76, 95], it is unclear how users explore and engage with LBSN services that can capture the *offline* and *online* experiences of the users. Do user engagement patterns in LBSNs present their own distinct patterns, or coincide with the human life course [25, 45]? That is, a person goes exploring in an “adolescent” phase and then becomes more stable by “settling down” later on. (2) Existing studies on churn prediction either focus on one user type such as newcomers [23, 86, 89], or one indicative feature set such as temporal [68, 80, 98], linguistic [4, 64, 76], or social feature sets [22, 67]. Hence, the effect of each feature and combinations of different features are not yet well understood. In particular, the churning of users who are significantly more active (i.e., post more reviews) than average users has received less attention from the research community. Thus, in this work, we focus on highly active producer-type users while defining the scope of the user engagement to the user behavior of writing a review on the platform for further analysis.

To overcome the limitations of the first class, we conduct a quantitative study to understand user engagement patterns in LBSNs better. For that, we employ two large-scale datasets of LBSNs: Yelp and Foursquare. To address the second class of limitations, we first characterize user types according to their contribution levels to the services and then we further analyze the distinct engagement patterns that producer-type users with significant contributions manifest themselves in various aspects in the datasets. To this end, we use the LBSN datasets to answer the following research questions:

- Q₁ How do highly active producer-type users engage in the services of LBSNs in terms of geographical exploration?
- Q₂ How do engagement patterns of highly active producer-type users manifest themselves in various aspects?
- Q₃ To what extent can we predict the churning of users with significant contributions within a given period of time?

1.1 Highlights of This Work.

We present a large-scale quantitative exploration of patterns of user engagement in LBSNs. We employ a dataset from Yelp to analyze user engagement. Next we include a dataset drawn from Foursquare for making our results generalizable to LBSNs based on user reviews. In Section 3, we describe the details of these two datasets. In Section 4, we present an analytical framework used throughout this work. The focus of this study is on user engagement over the whole lifespan of users and on the precise prediction of churning users. Thus, we further characterize users who produce UGC into two *long-term producers* and *ordinary producers* to analyze their engagement according to the level of contributions to the sites. We refer to those who made at least 50 reviews as *long-term producers* so that we can study users' engagement for sufficiently long periods of their lifespans. Then, the remaining users, i.e., the ones who wrote less than 50 reviews, are labeled as *ordinary producers*. This characterization facilitates our analyses on engagement, and churn prediction of long-term producers who write

significantly more reviews (write 18 times more reviews in Yelp and 20 times more in Foursquare as it can be derived from Table 1) than ordinary producers in LBSNs. We then describe the churn prediction problem in detail. To answer Q_1 , we examine how long-term producers engage with LBSNs in terms of geographical change. After that, we explore the differences in diverse aspects of long-term producers to answer Q_2 . Lastly, we formulate a prediction task to answer Q_3 .

To answer Q_1 , we start with the assumption that users in LBSNs would be more likely to explore geographically and then become less adventurous with age. To understand user engagement patterns for sufficiently long periods of their lifespans, we employ long-term producers. We first find that users' average radii and moving distances converge in a short time and are stable over their *lifecycle*, as defined in [21]. We then discover that users, in contrast to our initial assumption, continuously seek out different venues in new locations. We show that users return to the vicinity of previously reviewed venues from 10% to 40%, which means users tend to visit different venues with chances of 60–90%. These results can give insights for site maintainers to offer personalized venue recommendations by considering users' average radius and moving distances as well as their geographical engagement patterns.

For answering Q_2 and establishing principles to be used in prediction tasks, we examine the behavioral differences between churners and stayers among long-term producers from four aspects: (1) *geographic*, (2) *venue-specific*, (3) *social*, and (4) *linguistic*. Interestingly, behavioral differences between long-term producers who churn or stay are significant with their *first 10 reviews*. We find that long-term producers who stay consistently travel more to different locations and try more diverse categories of venues than those who churn. Besides, we discover that churned friends have more influence on long-term producers than on ordinary producers. Churning rates of long-term producers change more than twice than those of ordinary producers as the proportion of churning friends increases.

To answer Q_3 , we formulate a churn prediction task to distinguish churning users from staying users. It is crucial to identify potential churners from the group of long-term producers before they decide to leave the community because this group produces about 40% of reviews on Yelp and 20% of tips on Foursquare according to the analyzed datasets. We demonstrate that the insights established in this work enable us to predict whether a user will stop contributing to the community in the future. Our model based on Logistic Regression (LR) using all derived features achieves 0.768 AUC¹ (711 AUC), which outperforms all baseline models by up to 56.4% (43.3%) in AUC in Yelp (in Foursquare). After that, we explore to what extent we can further improve the performance of predicting churners by adopting a deep learning approach. Our best model achieves even higher performance of 0.882 AUC in Yelp and 0.799 AUC in Foursquare.

In summary, the main contributions of this work include:

- (1) We show that users constantly wander around diverse offline places, contrasting to the human life course assumption.
- (2) We find that the behavioral differences between churners and stayers are significant and that various factors show these differences with users' first 10 reviews.
- (3) We demonstrate the effectiveness of our observations by significantly improving the performance over all of the baseline LR models on the churn prediction task. Based on our proposed features, we employ a deep learning model and achieve even higher performance in predicting potential churners in LBSNs.

The rest of the paper is organized as follows. Section 2 summarizes the related work and Section 3 describes the analyzed datasets. Section 4 presents the designed analytical framework. In Section 5, we study users' geographical engagement. After that, in Section 6, we examine users' behaviors based on their reviews. In Section 7, we predict churning users. After that, Section 8 discusses the limitations of our approach and Section 9 describes the potential implications of this work. Finally, we conclude the paper in Section 10.

¹<https://developers.google.com/machine-learning/crash-course/classification/roc-and-auc>

2 RELATED WORK

We review the literature on user engagement, human mobility and revisitation patterns, and the growth and evolution of social networks as well as urban computing and LBSNs.

2.1 User Engagement

Researchers have investigated user engagement in various online settings such as mobile applications (apps) [53, 86], online communities [21, 76] as well as patterns [41, 61, 62] and motivations [47, 70] of user participation. In the literature, user engagement is formally defined as “the quality of the user experience that emphasizes the positive aspects of interacting with an online application and, in particular, the desire to use that application longer and repeatedly [46].” In addition, much work has focused on indicative features such as temporal [68, 80, 98], linguistic [4, 21, 64, 76], and social effects [22, 67] to improve user engagement or prevent users from churning. For instance, Danescu-Niculescu-Mizil et al. examined user engagement patterns from the language aspect and confirmed that users’ language usage becomes more inflexible with a community over time [21]. Amiri et al. used linguistic features from the tweets to predict the churning tweets for a specific brand of one telecommunication company among several others [4]. Instead of user churning from a single social network. Rashid et al. highlighted the similarity and acceptance of users with the rest of the community group are key factors for user participation on the platform [70]. Besides, Yang et al. used the network properties such as user’s network density and size with the user’s daily activities to predict churning users [86]. Lin et al. showed the primary intent of joining as the reason for multiple lives and proposed the method to predict the number of lives of the users [53]. Also, Mathur et al. modeled user engagement using contextual factors derived from smartphone usage and its embedded sensors [59]. Although Yang et al. [90] incorporated geographical information to improve the performance of the prediction of churning migrants from an urban area, it is still unknown how users’ location histories relate to their engagement patterns in LBSNs. Thus, in this work, we analyze user engagement patterns concerning geographical exploration using two large-scale datasets of LBSNs. We further study the geographical influences as well as venue-specific, social, and linguistic features on user engagement according to user types.

2.2 Human Mobility and Revisitation Patterns

Many prior works studied to reveal human mobility patterns [15, 16, 27, 35, 54–56, 93]. Specifically, Gonzalez et al. using mobile phone data found that human mobility displays significant regularity because they return to a few highly frequented locations such as home or work [27]. Cho et al. demonstrated that human movement patterns are periodic both spatially and temporally as well as highly influenced by social network ties when the movement is long-distance [15]. Moreover, Choi et al. conducted a field study to show the significant association between geographical exploration and a user’s information seeking behavior [16]. Lu et al. characterized the lifecycle of POIs and developed a framework to predict the life status of POIs in a given time slot [56]. Meanwhile, there are studies that investigate revisitation patterns in online [2, 36, 66, 69] and offline contexts [11]. Obendorf et al. showed that users’ navigation strategies on web pages differ dramatically and are largely affected by their habits and type of a site visited. The authors categorized users’ revisitation patterns three-fold according to heuristically defined time as follows: short-term (within an hour), medium-term (within a week), and long-term (longer than a week) re-visits [66]. Adar et al. conducted a large-scale analysis of Web interaction logs of about 612K users and characterized four fundamental revisitation patterns (e.g., Fast, Medium, Slow, and Hybrid groups) by clustering the revisitation curves they proposed [2]. This work has been extended to smartphone app usage and human mobility in urban spaces. Jones et al. found that the revisitation behaviors of smartphone users on a macro-level resembles those of web browsing on desktops [36]. In the contexts of urban areas, Cao et al. analyzed the physical revisitations of individuals and compared both similarities and differences of online and offline revisitation patterns [11]. On the other hand, our work focuses on the reviewing behavior of a user which reveals

one aspect of human mobility and other various factors manifested by users in LBSNs. We examine whether users show revisitation patterns to previously visited locations or show exploratory patterns to new locations while relating our findings to a practical application such as churn prediction. Note that we consider re-visiting a vicinity of previously visited locations instead of visiting the exact same locations again by writing a review since the data shows inconsistent patterns: users in Yelp are highly unlikely to write a review at the same venue again (0.04%), whereas users in Foursquare on average write two tips on the same location. To the best of our knowledge, this is the first work to investigate individual mobility concerning writing reviews on POIs.

2.3 Growth and Evolution of Social Networks

Studies on the growth and evolution of social networks are well-recognized research topics [7, 33, 38, 42, 50, 81]. Kossinets and Watts analyzed an evolving social network using email interactions among students over regular semesters. The authors found that although local connections between individuals evolve over time, the overall network structure remains stable [42]. Backstrom et al. found dense communities which have more closed triads grow less [7]. Leskovec et al. examined the microscopic process of the evolution of social networks [50]. In contrast to a global social network, an ego network is a personal social network consisting of two components: an *ego* and *alters* [6, 44]. The ego is a single user, while alters are directly connected neighbors to the user. Kikas et al. studied ego networks on Skype and observed that bursty peaks of contact additions tend to appear shortly after user account creation [40]. Aiello and Barbieri analyzed the temporal evolution of ego networks extracted from Flickr and Tumblr and found that users tend to build most of their ego networks in the early stages of their life [3]. Our work is relevant to those studies since users' churning behaviors can be influenced by their groups/communities (*i.e.*, social network properties) as well as their underlying social networks (*i.e.*, ego networks) as shown in prior works [22, 67]. In this work, we further extend those previous works by combining various ubiquitous data sources and significantly improve the performance of churn prediction tasks over a model built on social features.

2.4 Urban Computing and LBSNs

A new term, urban computing, has gained a significant attention which is largely attributed to the huge volume of data generated in cities [9, 12, 19, 20, 57, 73, 74, 77, 83]. These data-driven studies utilize different data sources ranging from GPS and mobile phone data to social media activity data. For example, Cranshaw et al. employed a GPS location tracking app to infer users' social relations from their physical locations [19]. Xu et al. utilized mobile phone and travel survey datasets to represent temporal modes in human trajectory as well as showed the correlation between popular temporal modes and user's occupations [83]. In order to improve the recommendation performance, prior works on POI [9] and event recommendation [20, 57] use users' activity data as well as their social and geographical information. In addition, common practices to collect data involves self-reporting through surveys or collecting data through some observational methods [32]. Tu et al. tackled the cold-start problem of the personalized location recommendation by learning user interest and location features from app usage data [77]. In the contexts of LBSNs, Chen et al. investigated user behaviors of cross-site linking according to their privacy concerns [14]. D'Silva et al. investigated influential factors that cause the failure of retail businesses and developed predictive models to foretell business survival using Foursquare check-ins and transport data [24]. Furthermore, Yang et al. [87] proposed a hypergraph embedding approach designed for LBSNs data and improved the performance of friendship and location prediction tasks. The hypergraph includes user-user edges (*i.e.*, friendships) and user-time-POI-semantic hyperedges (*i.e.*, check-ins). Xu et al. designed a deep learning pipeline for fine-grained Location Recognition and Linking and then showed the effectiveness of their framework on Twitter data [82]. However, no prominent work has studied the churn prediction problem in the contexts of LBSNs.

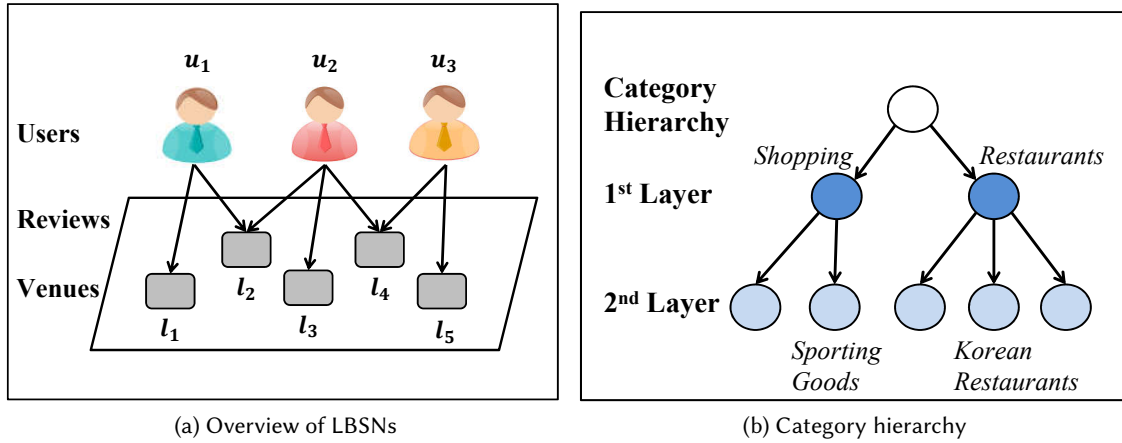


Fig. 1. Overview and a category hierarchy of location-based social networks.

3 DATASETS

In this Section we describe the two datasets we used in this work. We employ data from Yelp as our primary dataset to analyze user engagement. Then we include data from Foursquare to make our results more generalizable to LBSNs. The Yelp dataset is publicly available² and spans from July 2004 to December 2017 in four English-speaking countries. It is composed of over 1.3 million users, and around 5.3 million reviews of 175 thousands of businesses³. The Foursquare dataset, collected by Chen et al. [13], spans from October 2008 to February 2016 from around the world. It includes over 62 million users and 19 million tips from 13 million venues.

Figure 1 describes the key components of LBSNs: (1) *users*, (2) *venues*, (3) *reviews*, (4) *check-ins*, (5) *location history*, and (6) *category hierarchy*. Each user has basic profile information such as name, ID, friend list, and profile photos. In addition, when users visit a venue (e.g., a shopping mall), they can leave a review and a rating for the venue. If users mark the venue, it is known as a check-in to the venue. For example, as shown in Figure 1a, u_1 writes reviews of two venues which are l_1 and l_2 . For each review, the location can be extracted from the venue that the review has been written about. We can construct the whole location history of users using their reviews.

In contrast to other works [18, 92] which use users' check-ins to build their location histories, we use reviews to extract location histories of users for the following two reasons: (1) users often check-in to venues without physically visiting the locations as Zhang *et al.* found that nearly 75% of all Foursquare check-ins do not match with the real mobility of users [96]; (2) writing a review usually indicates that a user performed the relevant activities like shopping or dining at the specified venue [8]. In LBSNs, venues are grouped into pre-defined categories and Figure 1b depicts the category hierarchy in LBSNs. For example, the category "Restaurants" includes "Korean Restaurants", "American Restaurants", and so on. Although there are three or four layers of category hierarchies in LBSNs, we use two layers from the root since more than half of the venues do not contain the third layer's categories.

²<https://www.yelp.com/dataset>

³A review and business are called a tip in Yelp and venue in Foursquare. Hereafter we use these terms interchangeably. We restrict posts to reviews and use posts to refer to reviews.

Table 1. Descriptive statistics for datasets.

	Yelp	Foursquare
Number of total producers	132,291	1,209,210
Number of total reviews	1,558,344	7,650,575
Number of 1st layer categories	22	9
Number of 2nd layer categories	598	411
Number of long-term producers	4,730	15,403
Number of reviews of long-term producers	621,373	1,524,843
Average number of reviews	131.4	99.0
Average number of categories	41.2	38.6

4 ANALYTICAL FRAMEWORK

In this Section, we first characterize two types of users according to their contributions (Section 4.1) and after that we describe the churn prediction problem (Section 4.2).

4.1 Characterization of User Types

Considering that LBSNs largely rely on UGC, we focus on users who create UGC and define user engagement as their reviewing behaviors in our study. The focus of this study is on user engagement over the whole lifespan of users as well as on the precise prediction of churning users. Hence, we further categorize producer-type users into two groups according to their levels of contributions to the sites. First, we refer to those who made at least 50 reviews as *long-term producers* so that we can study users' engagement for sufficiently long periods of their lifespans. Second, we refer to the remainder of the users who wrote less than 50 reviews as *ordinary producers*. We set the number of contributions as the threshold used in prior works [21, 76] to distinguish long-term producers and ordinary producers so that our findings on LBSNs can be directly comparable with the previous works.

We choose long-term producers as the main focus of our work to study user engagement patterns over time, following [21, 76]. We identify 4,730 long-term producers in Yelp⁴ (15,403 in Foursquare⁵) who first posted reviews before January 2014 (March 2012) and wrote at least 50 reviews up until December 2016 (February 2015) in order to give them enough time (2014–2016 in Yelp and 2012–2014 in Foursquare) to accumulate 50+ reviews. It takes 820 and 720 days to accumulate 50 reviews on average in Yelp and Foursquare, respectively. Long-term producers are good subjects to study user engagement since their histories of activities in the community provide enough information to exhibit certain patterns. Moreover, long-term producers write 40% (20%) of reviews written by users who first posted in the period of consideration in Yelp (in Foursquare). Table 1 describes our final datasets which are used in the rest of this work. Note that a huge influx of users started from the year 2010 in both datasets, making a majority of long-term producers (99% in Foursquare and 89% in Yelp) who first posted after January 2009.

⁴Since the Yelp dataset is a business-centered dataset, the dataset contains partial information of users' whole review histories. Hence, we include users who have more than 50% of their whole review histories in our analysis. Whereas in the Foursquare dataset, entire tips of each user are collected with the maximum limit of 500 tips. In the end, we only sift out 512 long-term producers who have more than 500 tips in Foursquare.

⁵Since the Foursquare dataset was crawled from around the world, we include users whose more than 80% of reviews are written in English for further analysis in terms of language aspects. We apply the same process on the Yelp dataset.




Fig. 2. CDF of time gaps between sequences of reviews of long-term producers.

4.2 Churn Prediction Problem

Churn indicates the rate of loss of customers from companies in areas such as telecommunications [34, 63] or credit card services [65] where customers make a contract with a company. Like other online communities [1, 23, 29, 39, 89], users in online communities and LBSNs are not limited to a subscription contract. Hence, if a user v is inactive for a substantial period, then user v is considered to be a churned user, or otherwise a stayer. However, defining an inactive period as churn criterion differs depending on the context of the applications. To set the appropriate churn criterion in this work, we calculate the time gap of inactivity between two consecutive reviews of each user v . Figure 2 shows that over 95% of time gaps of users never exceed one year. We further check the re-engagement patterns of users which can be segmented into multiple disjoint active periods as suggested in [53]. When we set the inactive period as one year, around 60% in Yelp and 80% in Foursquare of long-term producers have only one lifespan (*i.e.*, one active period). Thus, in this work, we adopt one year as the inactive period and analyze user engagement from a perspective of a single life rather than from a perspective of multiple lives.

4.2.1 Problem Statement. We aim to relate our analysis to users' future activity status since it is important to predict the users' future status in advance. Thus we further categorize users into two groups based on whether they eventually abandon the service or not. We first define a date (1 year before the last date of the datasets) as the Start-Of-the-Future (SOF), similarly to [76]. The SOF is January 2017 in Yelp and March 2015 in Foursquare. We then define *staying users* as those who post at least one review as of the SOF; we define *churning users* as those who stop writing reviews as of SOF. As a result of our characterization of users and taking into account churn criterion, we identify 3,081 staying users and 1,649 churning users in Yelp as well as 6,153 staying users and 9,250 churning users in Foursquare. The ratios of staying to churning users are 65.1% to 34.9% in Yelp and 39.9% to 60.1% in Foursquare. Note that these long-term producers, who are the main focus of our study, had written at least 50 reviews before the SOF.

4.2.2 Notation. Given an LBSN $G(V, E)$, vertices are users and edges represent relationships between users. Set V denotes the users and set E their relationships. Each directed edge $e_{ij} \in E$ represents that user $v_i \in V$ follows another user $v_j \in V$. We call user v_i a "follower" of user v_j . If user v_j also follows user v_i , those users are "friends" with each other and we connect them with a directed edge e_{ji} . For each review users make, it contains location

and time information. Using this information, we can keep track of users' location histories (*i.e.*, sequences of venues the users have visited) along with the actual time period. In the following analysis, we use the time period as either *window* W or *stage* S . Let user v make T reviews. Then the entire indexed sequence $1, \dots, T$ can be grouped together as non-overlapping consecutive windows w_i where $i \in [1, T/\text{window size}]$ and stages s_j where $j \in [1, T/\text{number of stages}]$. The locations that user v has visited can be denoted as $L_v^w = [l_1, \dots, l_n]$, where location l_k consists of latitude and longitude.

5 GEOGRAPHICAL ENGAGEMENT OVER USER LIFECYCLE

In this section, we study user engagement from the perspective of users' *lifecycles* considering all the user's reviews. As in [21], we use the *life-stages* of a user to indicate the percentage of reviews the user has already written out of the total number of the reviews the user will write during her entire lifespan in the community. For example, a life-stage of 0%, *birth*, represents the time a user wrote her first review and a life-stage of 100%, *death*, represents the time a user leaves the community. In order to understand user engagement patterns over users' overall lifecycles, we employ long-term producers who have contributed more than 50 reviews. We investigate how users explore the real world in terms of their visited locations revealed by the history of their reviews.

User's average radius is determined in the early stage of their lifecycle. We use users' location histories to understand their geographical exploration. Given the entire location history, $L_v = [l_1, \dots, l_T]$, of user v , which is ordered by time and contains the latitude and longitude of the visited locations, we can compute user v 's *average radius* $r_g^v(t)$ using her trajectory up to t^{th} reviews.

$$r_g^v(t) = \frac{\sum_{i=1}^t |l_i - l_{CM}|}{t} \quad (1)$$

where $l_{CM} = \frac{\sum_{i=1}^t l_i}{t}$ is the center of mass of her trajectory. As studied in [27], users can be grouped into distinct groups according to their final $r_g(T)$. We group them into four so that each group contains approximately 25% of users within it (*e.g.*, 25% of users in Yelp and Foursquare are classified into a group with $r_g(T)$ less than 6 km and 80 km, respectively). Figures 3a and 3d show that the average radius of a user rapidly converges to $r_g(T)$ from the beginning of the user's lifespan. Moreover, a similar trend appears when we change the x-axis to the number of reviews from the life-stage as shown in Figures 3b and 3e. We show the average radii of users using their reviews are determined in the early stages of their lifetimes, which further confirms the result of prior work on human mobility using mobile phone data [27].

User's moving distance is constant over the user's lifecycle. Next, we examine the *average moving distance* of user v at each life-stage of her entire lifecycle. Given user v 's location history at each life-stage s , $L_v^s = [l_1, \dots, l_n]$, we calculate the average moving distance at each life-stage s as follows:

$$\text{Average moving distance} = \frac{\sum_i |l_i - l_{i-1}|}{|L_v^s|} \quad (2)$$

Figures 3c and 3f show the average moving distances of users according to their $r_g(T)$. The average moving distance, similar to the average radius, is determined in the early stages of the users' lifecycles. Besides, it remains consistent over the users' entire lifetime, meaning the users keep moving around geographically.

Will users settle down or keep exploring geographically diverse venues for reviewing? We examine whether user engagement patterns are analogous to the course of a human life [25, 45], *i.e.*, a person explores in a "adolescent" phase and then stabilizes by "settling down". To validate this question in terms of users' geographical locations, we start by investigating the human movement by measuring how likely user v is to travel close to the venues that she has written reviews about before. We analyze how often users return to previously visited venues and in the vicinity of them by writing reviews. Intuitively, we expect that users tend to explore diverse venues in geographically different locations in their earlier life-stages and then tend to explore less while sticking to their

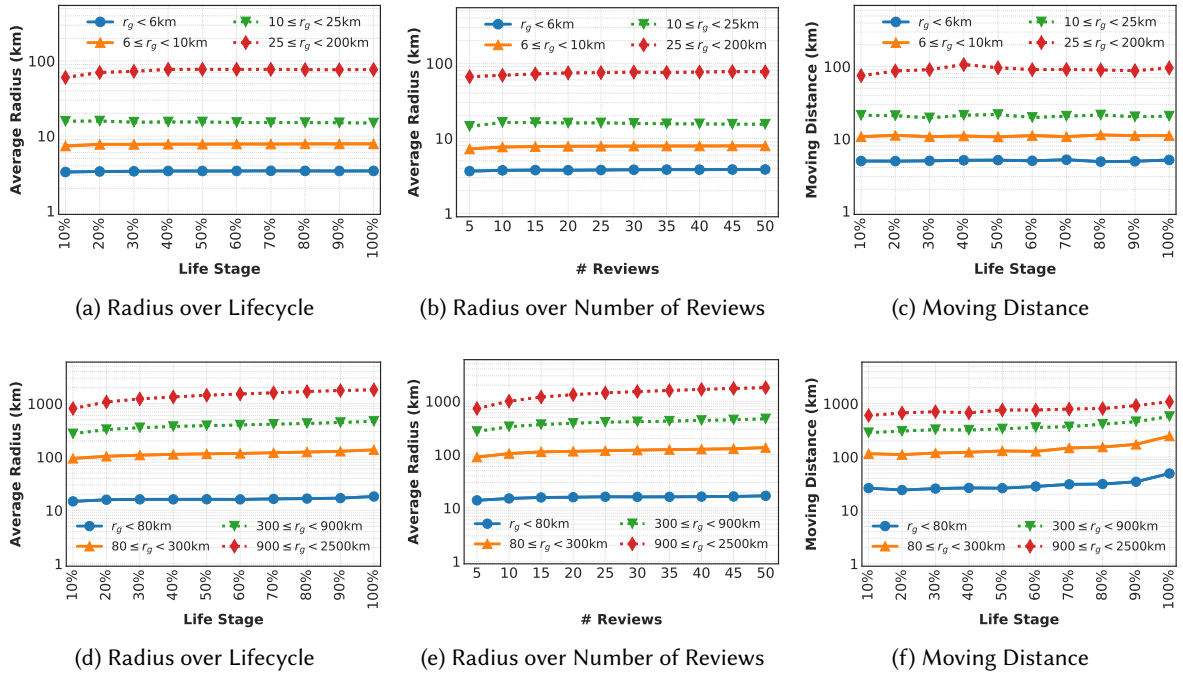


Fig. 3. Geographical engagement patterns. (a) The average radius of a user at each life-stage from an entire location history over her lifecycle. Users' final $r_g(T)$ can be grouped into 4 distinct groups and the average radius of users converges from the initial stages of their lives. (b) The average radius of a user for her first x reviews. It converges after the first few reviews. (c) The average moving distance of a user at each life-stage. Users' moving distances are stable in most of their life-stages. Standard-error intervals are depicted but very small. (Same trends of (a), (b), and (c) in Yelp hold for (e), (f), and (g) in Foursquare, respectively)

preferred places. More specifically, the likelihood of reviewing geographically different locations is quite high in the beginning stages of the users' lives, and then it decreases as users visit more and more places to review over their lifetimes in LBSNs. However, we will show below that this is not the case.

To quantify the propensity of users to explore diverse locations, we first define a counting function f to represent the number of occurrences of a user v reviewing the vicinity of the visited venues within radius d as follows:

$$f(L_v^{s_i}, L_v^{s_j}) = |\{l_s \in L_v^{s_i} : \text{distance}(l_s, l_r) \leq d \ \&\& \ l_r \in L_v^{s_j}\}| \quad (3)$$

where s_i is the current life-stage and s_j the previous life-stage. We then compute $P_{prev}(L_v^{s_i})$ which measures the probability that a user v in the current life-stage s_i revisits the vicinity of the venues for reviewing within radius d explored in the immediate life-stage s_{i-1} as follows:

$$P_{prev}(L_v^{s_i}) = \frac{f(L_v^{s_i}, L_v^{s_{i-1}})}{|L_v^{s_i}|} \quad (4)$$

$P_{prev}(L_v^{s_i})$ represents how likely the user v is willing to travel to various geographical locations in each stage of her life. For instance, $P_{prev}(L_v^{s_i})$ of 0% represents user v always writes reviews to different neighborhoods in her current life-stage s_i and $P_{prev}(L_v^{s_i})$ of 100% represents user v always return back to the vicinity of venues

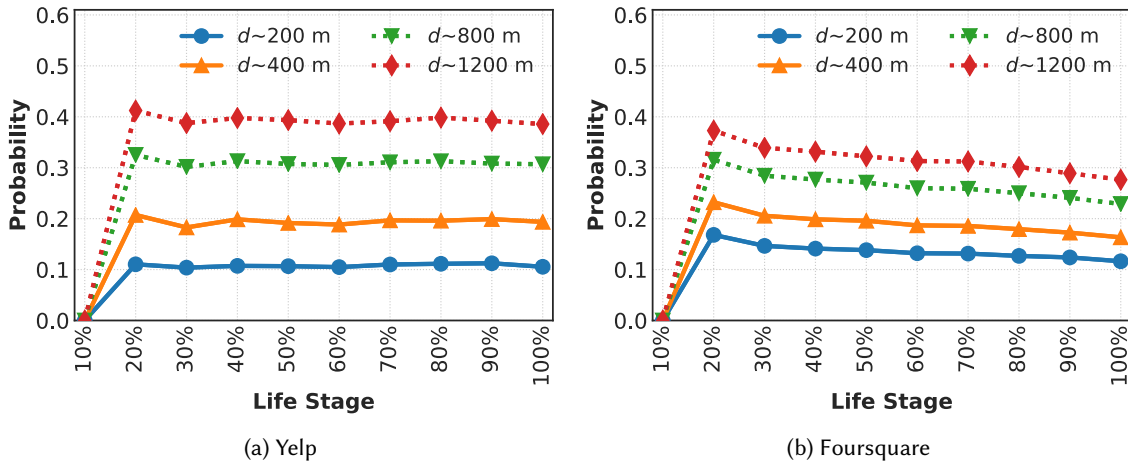


Fig. 4. The average probability of a user reviewing the vicinity of previously traveled locations in her immediate life-stage. Users keep exploring diverse venues in each life-stage over their lifecycle.

toured in her immediate life-stage s_{i-1} . We analyze users' geographical exploration using various threshold values of d defining the distance of the vicinity and observe similar trends. As shown in Figure 4, users in Yelp and Foursquare return to the vicinity of previously reviewed venues from 10% to 40% of the time. This result indicates that there is a 60–90% chance that users consistently post on different locations that they have not yet explored. Note that a radius of 400 meters is established as a standard for defining the size of a neighborhood which shares a similar functionality in the urban planning research community [60, 94].

We have observed that users are more likely to visit distinct locations when we only consider venues reviewed in their immediate life-stage s_{i-1} . However, people may return to the venues that they have reviewed in any life-stage s_j where $j \in [1, i-1]$. Hence, we calculate $P_{total}(L_v^{s_i})$ which measures the probability of the users in life-stage s_i revisiting the vicinity of all of the traveled venues for reviewing until the immediate life-stage s_{i-1} as follows:

$$P_{total}(L_v^{s_i}) = \frac{\sum_{j \in [1, i-1]} f(L_v^{s_i}, L_v^{s_j})}{|L_v^{s_i}|} \quad (5)$$

The more users accumulate reviews, the easier it is for them to write reviews on venues which are in the vicinity of previously visited venues. Surprisingly, as in Figure 5, the probability of visiting different venues converges to 40–70% in Yelp and 30–50% in Foursquare. In other words, there exist 30–60% (50–70%) of chances that users in Yelp (in Foursquare) keep exploring geographically different neighborhoods depending on the threshold distances d defining the vicinity. This result indicates that there is still a high chance that people travel to geographically distinct places even when we take into account all of the venues that they have posted reviews on so far. As a result, we validate users' geographical exploration throughout their lifecycle.

Section summary. The average radii and the average moving distances of users are settled soon after they start writing reviews on the site. However, users keep wandering around geographically diverse neighborhoods as they contribute more and more reviews in the site.

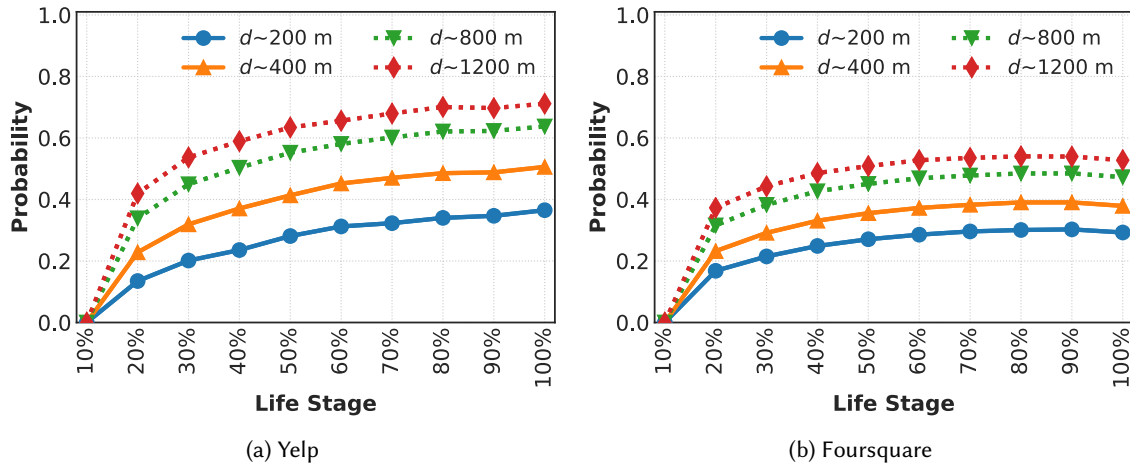


Fig. 5. The average probability of revisiting the vicinity of all previously traveled locations of users for reviewing. Regardless of various threshold values of determining the distance of the vicinity ($d = 200, 400, 800$, and 1200 m), the probability converges. This result shows that users keep exploring geographically diverse venues with chances of at least 30–60% in Yelp and 50–70% in Foursquare over their lifecycle.

6 ENGAGEMENT OF CHURNING AND STAYING USERS

After pointing out that users keep exploring diverse locations over their lifecycle in LBSNs. We now turn our attention to differences in user engagement between churners and stayers so that we can derive relevant feature sets to use in prediction tasks. In this section, to provide a holistic viewpoint to consider diverse aspects of ubiquitous data of LBSNs, we propose to investigate features related to location contexts such as geographical and venue-specific factors which LBSNs originally provide as well as social and linguistic factors as studied in previous works. Hence, we quantitatively study how churning behaviors manifest themselves among long-term producers using the following aspects: (1) *geographic*, (2) *venue-specific*, (3) *social*, and (4) *linguistic* aspects. To extract features for the prediction tasks, we take the initial 50 reviews rather than having $x\%$ of reviews because it is hard to know users' entire lifespan and what percent of their life has passed before their departure. Hence, in this section, we conduct our analyses using the initial 50 reviews of producers. After that, in Section 7, we show the effectiveness of the derived features in churn prediction in which we attempt to detect churning users early in their lifecycle using the initial x reviews.

6.1 Geographic Aspects

We delve into differences between churning and staying users manifested by geographical aspects. We choose four geographical features with which we examined the user engagement over the lifecycle in Section 5.

Average Radius. We use the average radius to investigate how this feature is related to the churning rates of long-term producers. Figures 6a and 6b show the decreasing trend of churning rates as the average radius increases. Each point in Figure 6a and 6b corresponds to the $r_g(T)$ values used in Section 5. In Yelp, the churning rate is 50% at the smallest $r_g(T)$ of 6 km. It becomes significantly reduced to 30% at $r_g(T)$ of 10 km and then is consistent onward. On the other hand, in Foursquare, the churning rate is 70% at the smallest $r_g(T)$ of 80 km. Then it continually reduces to 50% as the average radius increases.

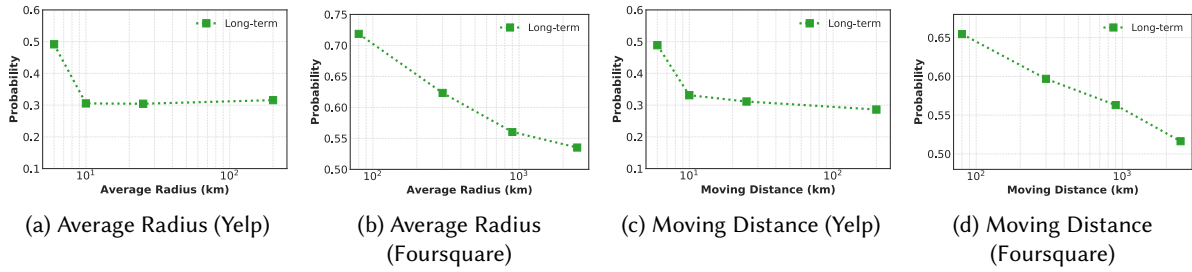


Fig. 6. Churning probability according to the average radius and moving distance.




Fig. 7. Revisiting probability according to P_{prev} and P_{total} .

Average Moving Distance. Similar to the average radius, the probability of churning rates is reduced according to the average moving distance as in Figures 6c and 6d. In Yelp, the trend of the churning probability of the average moving distance is shown to be almost identical to that of the average radius. In Foursquare, the churning rate declines from 65% to 50% as the average moving distance increases.

Revisiting probability in an immediate window. To further examine the churning behavior of users, we adopt the revisiting probability to the vicinity of venues in an immediate window based on written reviews of those users. Figures 7a and 7b show the average probability of writing reviews to the vicinity of venues that users visited in an immediate window. Staying users in Yelp are 4–6% and those in Foursquare are 2–3% less likely to write reviews than churning users from a neighborhood that they have already reviewed. This result indicates that both stayers and churners do not tend to return to the neighborhood of previously reviewed venues to write reviews again. Besides, stayers are relatively more likely to review new venues than churners.

Revisiting probability in all previous windows. Figures 7c and 7d shows the average probability of writing reviews to the vicinity of venues that users visited in all previous windows. Staying users in Yelp and Foursquare are around 4–6% on average less likely to write reviews than churning users from a neighborhood that they have already reviewed. This result further confirms two findings from the analysis of revisiting probability in an immediate window on human mobility based on reviewing behaviors of producer-type users.

6.2 Venue-specific Aspects

We employ several venue properties to study how user engagement relates to venues. Thus we use venue categories and the number of accumulated reviews written on a venue when a user wrote her review on it.

Venue categories. We employ second-layer categories (see Section 3 for detail) to analyze how venue properties affect user engagement. Figures 8a and 8b show the average number of unique second-layer categories in




Fig. 8. Average number of unique categories in each window.




Fig. 9. The number of accumulated reviews on a venue in each window.

each window. Stayers are more likely to visit diverse categories of venues than churners. In addition, the same result also holds for other metrics such as entropy and Gini index. Both metrics are based on the probability of categories C in each window w_i . The probability of a category c in a window w_i is computed as $p_c = \frac{1}{|w_i|} \sum_{k \in w_i} I(C_k = c)$. Then the entropy and Gini-index in each window w_i are defined as Eq. 6 and Eq. 7, respectively.

$$\text{Entropy} = -\sum_c p_c \log_2(p_c) \quad (6)$$

$$\text{Gini-index} = 1 - \sum_c p_c^2 \quad (7)$$

In sum, users write reviews on categories more evenly in a window. In addition, churners focus on fewer categories.

Accumulated Reviews on venues. We further investigate the venue properties using the number of reviews on a venue when a user visits it. We study whether the number of accumulated reviews on a venue affects users' decisions to visit the venue. As shown in Figures 9a and 9b, we find that the average number of accumulated reviews on a venue increases for both churning and staying users. Moreover, churners write reviews on venues with a fewer number of reviews accumulated than the venues on which stayers write reviews. Note that the

y-axis of Figures 9a and 9b are normalized using a log with a base of 2, whose values are used as an input to classifiers in Section 7.

6.3 Social Aspects

Using social network properties of LBSNs, we examine the social influence on users' churning behaviors. We compare the churning probability of long-term and all users to identify how much social influence affect churning behaviors of long-term producers who are the main focus of this work. Note that the employed data contains a static social network at the end of the inspection period. Hence, we do not have social network information of users when they first start writing reviews. However, based on the finding of the prior work that the ratio of added nodes to one's social network is considerably smaller after 20 days compared to the final ego-network size (Only nodes who have created links for at least 6 months are considered) [3], we could know that long-term producers' social networks may stabilize over time. Since it takes on average 820 days in Yelp and 720 days in Foursquare to accumulate 50 reviews, we assume that social networks of long-term producers already stabilized when they write around 50 reviews. Based on this assumption, we analyze the social aspects of user engagement and then conduct an experiment only on 50 reviews in Section 7.

Degree. We first examine the relationship between churning probability and the number of friends that a user has. Figures 10a and 10b show decreasing trends in churning rates as the number of friends increases. This tendency matches with our common sense that users with many friends in the service are less likely to leave. Note that we use the number of friends as a degree because the datasets do not provide follower relations but friend relations. In sum, we show that churning rates of all users are consistently higher than those of long-term producers. Furthermore, we confirm that the churning probability declines for all and long-term producers in proportion to the degree of the users.

Proportion of Churned Friends. Figures 11a and 11b show the churning rates of all and long-term producers according to the proportion of their churned friends, respectively. We find there are more churning users as they have more and more churned friends, as studied in [67]. In our study, we take all users as a reference point to examine the churning behaviors of long-term producers. The churning rates are very high of 80–100% for all users and 30–70% for long-term producers at the X-axis of 0%. These high churning rates are reasonable since the X-axis of 0% occurs when the user has no friends. Furthermore, we discover that long-term producers are more sensitive to their friends' churning. As in Figure 11a, churning rates of long-term producers increase significantly from 5% to 60% as the proportion of churned friends is raised from 20% to 80% in Yelp, which is almost twice as much change in the churning rates of all users. Also, Figure 11b shows a more significant change in churn rates of long-term producers in Foursquare. Thus, while the churning rates of all users are always higher than those of long-term producers, long-term producers are more sensitive to their friends' departure from the community.

6.4 Linguistic Aspects

We finally investigate how linguistic aspects affect users' churning behaviors using review text. We herein adopt language patterns that users use as a proxy to look into users' online interactions and engagement patterns in the community. We first take review lengths. Then, as studied in [17, 72], we take the frequency of pronouns to study the differences between churners and stayers.

Review length. We find that long-term producers in both Yelp and Foursquare write longer reviews while they accumulate more and more reviews on the site, as shown in Figures 12a and 13a. Moreover, in Yelp, the review length of churning users is significantly longer than that of staying users. In Foursquare, churning users write longer reviews in their first 10 reviews. However, the difference in review length between churning and staying users disappears after they accumulate more than 10 reviews. It is noteworthy that long-term producers in Yelp write much longer texts than the users in Foursquare. These differences between Yelp and Foursquare in




Fig. 10. Churning probability according to the number of friends.

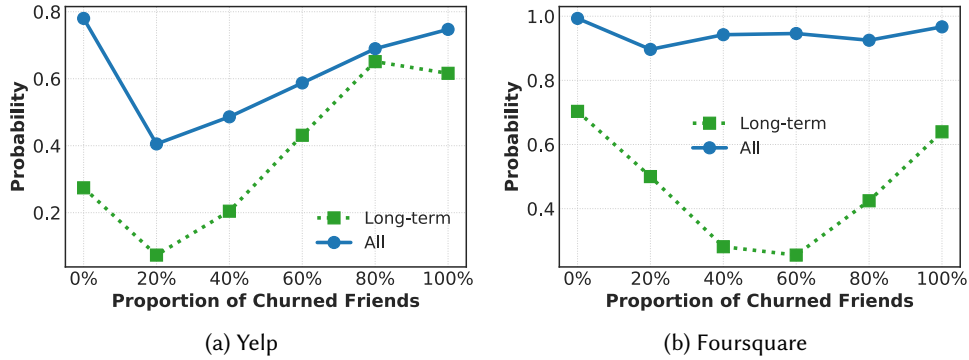


Fig. 11. Churning probability according to the proportion of churned friends.

terms of review length could be derived from the characteristics of each LBSN. Specifically, Yelp encourages long and detailed reviews on venues by displaying an elaborate and exemplary review to users. On the other hand, Foursquare promotes concise and brief tips on venues by asking a simple question (*e.g.*, What's good here?) with the limit of word count on tips.

Frequency of pronouns. Prior works [21, 76] studying linguistic features in online communities suggested that the decreasing frequency of first-person singular pronouns (*e.g.*, I, Me) can indicate the users' increasing identification with the community. In our study, however, Figures 12b and 13b show the opposite trend, *i.e.* the frequency of first-person singular pronouns manifest increasing patterns as users post more reviews on the site.⁶ If we follow the interpretation that the low frequency of first-person singular pronouns is associated with a higher level of affiliation with the community, it is hard to explain the increasing patterns of pronoun usage in LBSNs because it means newcomers who will churn have the highest affiliation with the community. Thus, we present a new perspective based on the linguistic theory that can better explain the observed language patterns in LBSNs.

⁶Although Figures 12b and 13b represent both first-person singular and first-person plural pronouns, first-person singular pronouns show a similar trend of Figures 12b and 13b.

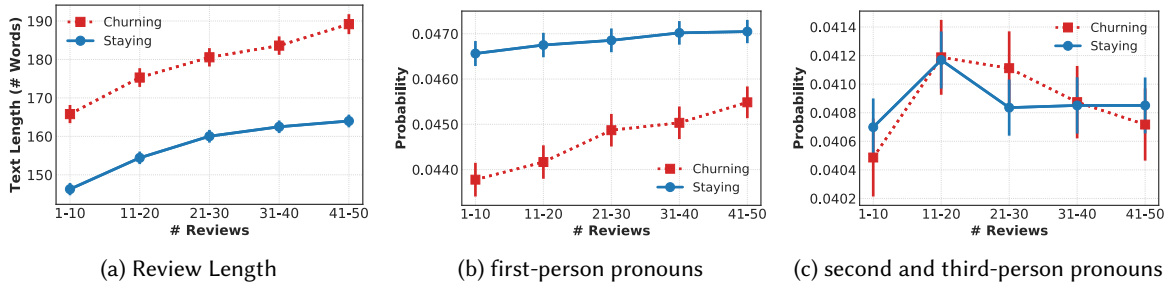


Fig. 12. Linguistic aspects in Yelp.

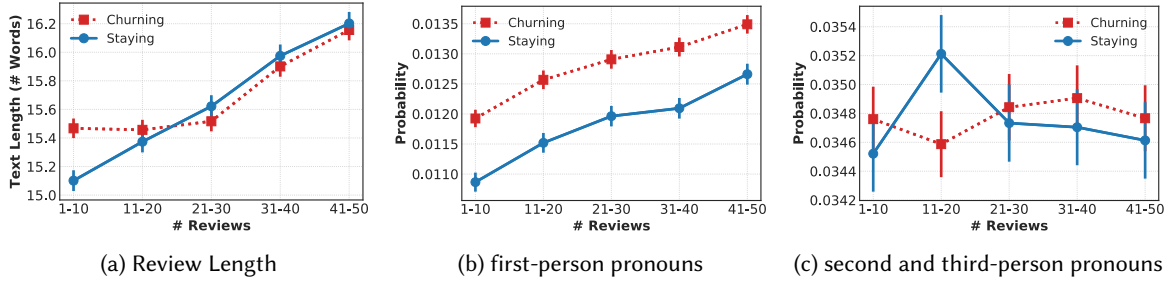


Fig. 13. Linguistic aspects in Foursquare.

In the literature, the frequencies of pronouns used by a user can be associated with the focus of the user [17, 37, 72]. For example, if a user often uses first-person pronouns (e.g., I, Me, We, Us), this indicates that the user's attention is on herself, friends, or family members within her group. On the other hand, the frequent usage of second and third-person pronouns (e.g., you, yours, they, theirs) can be associated with a user's attention on others who are not necessarily within her group. Users in LBSNs are likely to use more first-person pronouns as they contribute reviews to the community (see Figures 12b and 13b). In contrast, Figures 12c and 13c show that the usage of second and third-person pronouns represents consistent patterns over users' lifespans. This result may indicate that users focus more and more on people within their groups. Note that a small fraction of difference in pronoun frequencies can reveal meaningful behavioral differences in language patterns of users as in [76]. However, since the actual difference is quite small, linguistic feature sets extracted from frequencies of the pronouns are not as effective as other proposed features in the prediction task in Section 7 (see Table 4).

7 PREDICTING CHURNING USERS

Having established engagement patterns from various aspects, we turn now to study to what extent we can predict churning users from staying users using the identified engagement patterns. Predicting churners among a group of long-term producers has practical value for creating and maintaining online communities since the LBSNs primarily rely on the type of users who actively contribute 40% of all reviews. Furthermore, given that almost 70% of customers in subscription services will not come back once they stop the subscription [43], early detection of long-term contributors who are likely to leave the community is crucial to service maintainers by enabling them to use many strategies to re-involve the users in the services before they abandon those services.

7.1 Experimental Setup

We formulate a prediction task of detecting churning users using an initial k reviews. We then adopt the oversampling method to tackle the class imbalance of the datasets. Moreover, we vary the initial k reviews of users to examine the performance changes according to various first k reviews used for training classifiers.

7.1.1 Baseline. Prior work on user engagement found that a temporal feature, time gaps between reviews, is a powerful indicator to predict users' churning [68] and has often been used for training classifiers in recent works for user engagement in online communities [21, 76]. We build a benchmark using this temporal feature as a reference point and compare it with our proposed features to show the effectiveness of our proposed features on the classification task.

7.1.2 Proposed Features. We propose features based on the observations that we have reported in the previous sections. Table 2 describes four sets of features in detail. For all features except for social features, we use the aggregated values in each window with a size of 10 posts. Hence, we have five values for each feature. To take into account the temporal dynamics of features, we include the index of the window with the maximum and minimum value. Then we use each set of features to build the models for churn prediction task as follows.

- (\mathcal{F}_1) **Temporal feature:** Average time gaps between reviews.
- (\mathcal{F}_2) **Geographic feature:** Average radius, average moving distance, P_{prev} , and P_{total} .
- (\mathcal{F}_3) **Venue properties:** Unique category, entropy, gini-index, and average number of accumulated reviews on a venue.
- (\mathcal{F}_4) **Social feature:** Number of friends and churn rates of friends.
- (\mathcal{F}_5) **Linguistic feature:** Average frequency of first, second, and third-person pronouns as well as average number of words.

To further investigate the most prominent features based on our observations in Section 5 and Section 6 for learning models, we build models using the top-2 important features listed in Table 4. Also, we build models using the top-2 important features and two geographic features, one of the primary contributions in both Section 5 and Section 6. Finally, we construct the full model using all features and leave-one-out models from the full model.

- (\mathcal{F}_6) **Top2:** Average time gaps between reviews and average number of accumulated reviews on a venue.
- (\mathcal{F}_7) **Top2+Geo2:** Average time gaps between reviews, average number of accumulated reviews on a venue, average radius, and average moving distance.
- (\mathcal{F}_8) **All:** Combination of all features.
- ($\mathcal{F}_{9:15}$) **Leave-one-out:** Combination of all features without one feature set.

7.1.3 Methods for Evaluation. We use Logistic Regression (LR) with L2-regularization as a classifier to predict churners. We adopt the LR model which can provide us with highly interpretable information on our derived features since the primary goal of our study is to identify and validate important features from the observations we made in Section 5 and 6. Thus, we train LR for all proposed combinations of features (*i.e.*, $\mathcal{F}_1:\mathcal{F}_{15}$). Furthermore, inspired by the recent advancements in the deep learning approach for sequential data [49, 86], we adopt the Long Short-Term Memory (LSTM) recurrent neural network (RNN) widely used for time-series analysis [28, 51]. After that, to explore to what extent we can further enhance the performance of the churn prediction task, we train LSTM using all features (\mathcal{F}_8) which shows the best performance among all LR models. In experiments for LSTM models, we leverage Adam [48] as the optimizer and implement them with TensorFlow architecture. We set the batch size and learning rate to 32 and 0.001, respectively. Besides, we adopt the Glorot initialization [26] and early stopping [71] in the training process. The dropout probability [75] is set to 0.1 at the last LSTM layer. For both LR and LSTM models, we optimize the hyperparameters with the grid search strategy.

Table 2. List of Proposed Features of a User v

Geographic Features	
Average Radius	Average distance of v up to t^{th} review from her center of mass l_{CM} . i.e., $\frac{1}{t} \sum_{i=1}^t l_i - l_{CM} $ where $l_{CM} = \frac{1}{t} \sum_{i=1}^t l_i$.
Moving Distance	Average distance that v moves in window w . i.e., $\frac{1}{ L_v^w } \sum_i l_i - l_{i-1} $.
P_{prev}	Probability $P_{prev}(L_v^{w_i})$ that v visits venues within radius d of venues reviewed in an immediate window w_{i-1} in current window w_i , i.e., $P_{prev}(L_v^{s_i}) = \frac{f(L_v^{s_i}, L_v^{s_{i-1}})}{ L_v^{s_i} }$
P_{total}	Probability $P_{total}(L_v^{w_i})$ of v in current window w_i reviewing the vicinity of all of the venues v has travelled for all windows w_j where $j \in [1, i-1]$, i.e., $P_{total}(L_v^{s_i}) = \frac{\sum_{j \in [1, i-1]} f(L_v^{s_i}, L_v^{s_j})}{ L_v^{s_i} }$
Venue Properties	
Unique Category	Average number of unique second-layer categories.
Entropy	Category diversity based on the probability of categories C in each window w_i . i.e., $-\sum_c p_c \log_2(p_c)$.
Gini-index	Category diversity based on the probability of categories C in each window w_i . i.e., $1 - \sum_c p_c^2$ where probability of a category c is computed as $p_c = \frac{1}{ w_i } \sum_{k \in w_i} I(C_k = c)$.
# Accu. Reviews	Average number of accumulated reviews on a venue when v write a review on the venue.
Social Features	
Degree	Number of friends v has.
% Churned Friends	Percentage of churned friends of v .
Linguistic Features	
Review Length	Average number of words in v 's reviews.
1st person	Average frequency of first-person pronouns used by v .
2nd person	Average frequency of second-person pronouns used by v .
3rd person	Average frequency of third-person pronouns used by v .

7.1.4 Evaluation Protocols. We define the task to predict whether a user will churn or stay after their 50th review. To distinguish churning producers from staying producers, we extract features based on users' first k reviews where $k = 10, 20, 30, 40, 50$. For social features, we conduct experiments only with the case of 50 reviews due to the limitation of static social network data. Since the proportions of churners and stayers are imbalanced, we balance the proportion of two classes by oversampling the minority class (i.e., churners) [10]. In order to overcome the potential bias in our sampled datasets and to obtain the generalizability of our results, we conduct the experiments over 20 randomly sampled datasets. We determine 90% of users as training/evaluation sets and the remaining 10% of users as a test set, respectively. Then we use the area under the receiver operating characteristic curve (AUC) to evaluate the performance of models. AUC is a widely used measure to assess the performances of classifiers in imbalanced data [10, 53].

Table 3. Results of predicting whether a user will leave the community in the future. The prediction performances of LR models using each set of the proposed features and LSTM model using all features are presented for both Yelp and Foursquare. -Temporal represents LR model using all features except temporal feature, and the same holds for the rest of the features.

	Yelp					Foursquare				
Feature	$k = 10$	$k = 20$	$k = 30$	$k = 40$	$k = 50$	$k = 10$	$k = 20$	$k = 30$	$k = 40$	$k = 50$
Temporal	0.659	0.692	0.696	0.700	0.700	0.634	0.668	0.681	0.683	0.683
Geographic	0.589	0.605	0.604	0.609	0.609	0.577	0.594	0.604	0.606	0.605
Venue	0.616	0.622	0.640	0.659	0.661	0.566	0.572	0.587	0.598	0.606
Social	-	-	-	-	0.714	-	-	-	-	0.619
Linguistic	0.583	0.599	0.602	0.607	0.613	0.520	0.540	0.545	0.545	0.547
Top2	0.675	0.704	0.708	0.713	0.713	0.624	0.664	0.680	0.682	0.683
Top2+Geo2	0.681	0.710	0.713	0.722	0.723	0.630	0.668	0.686	0.688	0.689
All	0.687	0.715	0.720	0.729	0.768	0.633	0.671	0.689	0.692	0.711
-Temporal	0.643	0.659	0.669	0.684	0.736	0.594	0.615	0.628	0.634	0.661
-Geographic	0.680	0.706	0.716	0.722	0.767	0.624	0.667	0.683	0.685	0.708
-Venue	0.673	0.709	0.713	0.721	0.768	0.630	0.670	0.688	0.692	0.710
-Social	-	-	-	-	0.728	-	-	-	-	0.691
-Linguistic	0.680	0.712	0.717	0.723	0.763	0.634	0.670	0.686	0.689	0.710
-Top2	0.611	0.637	0.638	0.647	0.721	0.589	0.609	0.619	0.622	0.654
-(Top2+Geo2)	0.586	0.630	0.634	0.642	0.722	0.541	0.566	0.577	0.577	0.643
Stacked LSTMs	0.735	0.844	0.847	0.858	0.882	0.664	0.751	0.770	0.773	0.799

7.2 Evaluation on LR

The performances for predicting whether a user will depart the community in the future are shown according to the number of first k reviews used for training and testing the LR models (see Table 3). LR models trained with Top2 and Top2+Geo2 features show improvement over the benchmark model, which indicates that the number of accumulated reviews and two geographical features provide additional information onto the benchmark model trained with a strong temporal feature. In addition, our full LR model which utilizes all features outperforms all other models, by achieving 0.768 AUC in Yelp (0.711 AUC in Foursquare). Note that a random baseline will show 0.50 AUC. The full LR model also significantly improves the performance of the benchmark which uses strong indicator, temporal feature, by 9.7% (4.2%) in AUC and other baseline models by up to 56.4% (43.4%) in AUC in Yelp (in Foursquare). The differences between the benchmark and full LR model for both prediction tasks in Yelp and Foursquare are statistically significant according to the Wilcoxon signed-rank test ($p < 0.001$). This result validates the effectiveness of our suggested features in distinguishing churning users.

Furthermore, we conduct experiments with LR models using leave-one-out feature sets (*i.e.*, rows of -Temporal, -Geographic, -Venue, -Social, -Linguistic, -Top2, and -(Top2+Geo2)) to scrutinize how performance decreases with cutting part of the component. When using the first $k = 10, 20, 30, 40$ reviews, the performance of all leave-one-out LR models decreases. However, when using the first $k = 50$ reviews, the addition of Geographic and Venue features to LR models does not lead to performance improvement. It seems that geographic and venue-specific features do not provide much information by including more reviews since those features are consistent over time. On the other hand, adding social features when using the first $k = 50$ reviews significantly

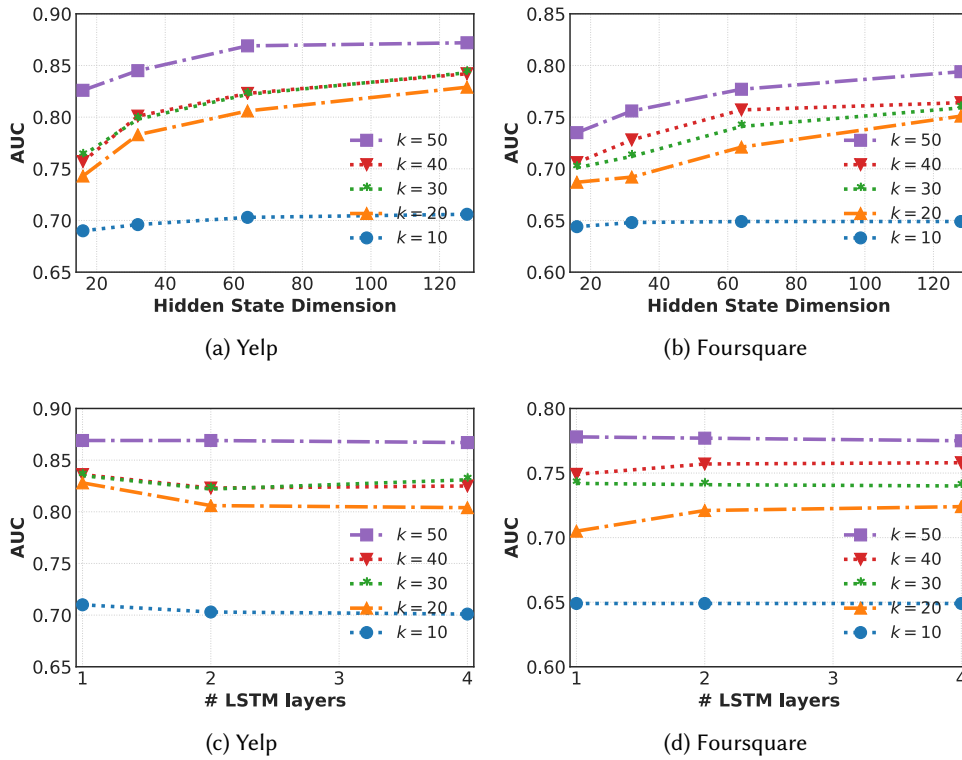


Fig. 14. Parameter sensitivity study on the performance of Stacked LSTMs.

improves the overall performance of LR models. This result indicates that using all possible features is essential for an improved performance as the full LR model performs the best.

7.3 Evaluation on Stacked LSTMs

Based on the result of LR models, we conduct further experiments to investigate to what extent we can improve the overall performance leveraging the recent advancement of a deep learning approach. For that, we utilize Stacked LSTM recurrent neural networks [28] using all features to compare with the full LR model. Table 3 also lists the evaluation results of Stacked LSTM models according to the first $k = 10, 20, 30, 40, 50$ reviews with both Yelp and Foursquare datasets. We observe that Stacked LSTM outperforms the best performing LR model over all cases. For the Yelp dataset, Stacked LSTM improves 6.9–18.0% over the best LR model. Also, Stack LSTM consistently achieves the best performance by improving 4.8–12.4% over the best LR model in the Foursquare dataset. In the end, the best Stacked LSTMs achieve a high AUC of 0.882 in Yelp and 0.799 in Foursquare.

Furthermore, we conduct experiments to investigate parameter sensitivity. Stacked LSTM involves several parameters (*e.g.*, hidden state dimension, the number of stacked LSTM layers, batch size, dropout probability). To examine the robustness of the trained Stacked LSTM models, we investigate how the performance of Stacked LSTM in predicting churning users is affected by the different choices of parameters. Except for the tested parameter (*i.e.*, hidden state dimension and the number of stacked LSTM layers), we set other parameters to the default values as specified in § 7.1. Figure 14 shows the evaluation results of Stacked LSTMs by varying two

Table 4. Feature Importance: χ^2 Statistics

Rank	Yelp			Foursquare		
	χ^2	Feature	Feature Category	χ^2	Feature	Feature Category
1	16625.17	# Accu. Reviews	Venue	8301.76	Time Gap	Temporal
2	5533.68	Time Gap	Temporal	7801.16	# Accu. Reviews	Venue
3	76.07	% Churned Friends	Social	439.84	Average Radius	Geographic
4	34.92	Average Radius	Geographic	344.47	Moving Distance	Geographic
5	28.92	Moving Distance	Geographic	23.54	Unique Category	Venue
6	14.36	Degree	Social	9.46	P_{total}	Geographic
7	12.36	P_{prev}	Geographic	8.54	P_{prev}	Geographic
8	12.01	P_{total}	Geographic	8.77	Entropy	Venue
9	6.22	Review Length	Linguistic	8.11	% Churned Friends	Social
10	1.95	Unique Category	Venue	4.64	Review Length	Linguistic

parameters. First of all, we observe that the change in Stacked LSTMs' performance is minimal when $k = 10$ with both parameters. In addition, we observe that the increase in the performance saturates as the hidden state dimension reaches around 64, which demonstrates that the larger dimensionality does not always bring performance increase. On the other hand, we find that the number of LSTM layers have a relatively low impact on the performance of Stacked LSTMs. This result indicates that even with a single layer LSTM can achieve high performance in predicting producer-type users in LBSNs.

7.4 Performance Change Over Reviews

We further discuss the change in performance with respect to the number of initial k reviews used for training the models. The performance of all models increases as we use more reviews for training the models as shown in Table 3. For example, a full LR model using 50 reviews displays an additional 11.7% (12.3%) improvement in AUC over the full LR model using 10 reviews in Yelp (in Foursquare). Similarly, the performance of Stacked LSTMs based on $k = 50$ posts improves by 20.1% (20.4%) in AUC over the model based on $k = 10$ posts in Yelp (in Foursquare). This result indicates that considering temporal dynamics by acquiring more reviews is as essential as having informative features. Note that the Stacked LSTM model using the first 10 reviews also achieves 0.735 AUC in Yelp (0.664 AUC in Foursquare). It represents that the information in the earlier stage of a user's life provides enough predictive power to predict churning users accurately. This result is impressive since we can already make an accurate prediction of the future status of users from their first 10 posts.

7.5 Understanding Feature Importance

We finally investigate the feature importance of our proposed features. We calculate the χ^2 (Chi-square) statistic to evaluate the discriminative power of our proposed features [91]. Table 4 shows the top 10 most important features with χ^2 scores of Yelp and Foursquare datasets. Along with the temporal feature, features such as geographic and venue properties derived from offline context are top 5 important features. As we discussed the results of experiments with leave-one-out features, it may not be informative when we have full 50 reviews of producers. However, this result can indicate the impact of taking offline context into account for the prediction task in more common cases where we only have the limited information of users (e.g., 10 reviews or less). Other features such as linguistic, social features are also vital for constructing the powerful predictive model since the models using the combination of all features perform the best.

8 LIMITATIONS

There are limitations to this work due to the employed datasets. We discuss these limitations in this section.

- (1) **Data availability.** The Yelp dataset is a business-centered dataset, which contains whole review histories of businesses but only some portion of those of users. Although we have tried our best to remain the highest coverage for analysis, many Yelp users are sifted out during the preprocessing process. Thus, there may exist some bias in the analysis due to data availability. However, in the Foursquare dataset, we capture the whole review histories of users. Moreover, our analysis shows consistent patterns for both Yelp and Foursquare users in the geographical, venue-specific, and social aspects. The bias can be largely mitigated.
- (2) **Studied users.** We focus on the behaviors of long-term producers who contributed at least 50 reviews so that we had sufficient history per user to observe her trajectory in the community as well as in the real world. However, this user type makes up 3.6% (1.3%) of the user base in Yelp (Foursquare) during the studied period, which limits our study to a small portion of users in LBSNs. Furthermore, our study excluded consumer-type users since they do not leave any logs to analyze in our study.
- (3) **Engagement Diversity.** LBSNs enable users to interact with the services in many ways. For example, users can search through the sites to find the next destination to dine, read reviews of local restaurants in unexplored areas to decide their visits to the venue, and interact with other users through up-voting their reviews or by following them. In our study, we could not consider diverse aspects of user engagement such as reading reviews and interacting with other users, since the employed datasets do not contain such information. However, by narrowing down the scope of the user engagement of producer-type users to the activity of writing a review, we present interesting findings in this work, which leads us to useful implications as we discuss in Section 9.

9 DISCUSSION

We use this section to summarize our findings in users' geographical exploration patterns and user engagement in four different aspects. Note that here we present the correlation between user behaviors and examined features, not causal relationships. After that, we discuss the potential applications of our findings.

The following six points can summarize our quantitative study on LBSNs:

- (\mathcal{P}_1) The average radii and moving distances of users are determined within 5–10 reviews (§5).
- (\mathcal{P}_2) Users consistently write reviews on different locations at least 50% of all reviews for each life-stage (§5).
- (\mathcal{P}_3) Staying users are more likely to explore diverse locations than churning users (§6.1).
- (\mathcal{P}_4) Staying users are more likely to write reviews on venues of diverse categories and with more reviews accumulated (§6.2).
- (\mathcal{P}_5) The probability of churning increases for users with a higher percentage of churning friends (§6.3).
- (\mathcal{P}_6) Churning users, in Yelp, use less first-person pronouns and write longer reviews. In Foursquare, on the other hand, churning users use more first-person pronouns and reviews of approximately the same length as the staying users (§6.4).

Based on the findings on users' geographical exploration patterns, we first confirm the previous studies on human mobility that human movement patterns are periodic and regularized [15, 27]. We also observe that users keep reviewing diverse categories and different locations of venues in contrast to the human life course theory [21, 25, 45], where a person explores in a “adolescent” phase and then stabilizes by “settling down”. Our finding is in accordance with the prior work on users' community seeking behaviors in Reddit [76]. Furthermore, all of our discoveries, including various aspects of engagement patterns of users, have implications for site maintainers to increase user engagement in LBSNs. To increase the engagement levels of users, incentive mechanisms like gamification [30] can be employed to encourage engagement. For example, for those users

who are at risk of departing, site maintainers can provide some incentives such as rewards and badges or can recommend different venues that users have never reviewed before to re-engage them.

Our findings can be utilized in the following ways:

- (1) Based on (\mathcal{P}_1), early recognition of user geographical exploration patterns enable the site maintainers to provide a more personalized user experience. Since the average radii and moving distances of users are determined within 5–10 reviews and are stable over their lifespan, one can recommend venues located within a user’s average radius and moving distance from her center of mass. For example, for those who have a small average radius and moving distance less than 6km, one can suggest nearby venues to the users. On the other hand, for those who have a large average radius and moving distance greater than 100km, one can even recommend venues located in another city.
- (2) Based on (\mathcal{P}_2) and (\mathcal{P}_4), since users tend to write reviews on unexplored geographical locations and diverse categories of venues, we can recommend a venue with a new category and location in unexplored neighborhoods that the user has not yet visited for reviewing. For example, if a user has some reviews on a category of “Mexican Restaurant” in one neighborhood, one can recommend a “Chinese Restaurant” located in another location to encourage the user to explore and to increase her engagement with the services.
- (3) Based on (\mathcal{P}_4), staying users are more likely to write reviews on popular venues (*i.e.*, the high number of accumulated reviews). It seems that users are more satisfied with popular venues. Hence, we can suggest popular venues to users for reviewing to increase their engagement on the services.
- (4) Based on (\mathcal{P}_3), (\mathcal{P}_4), (\mathcal{P}_5), and (\mathcal{P}_6), the powerful predictive model leveraging various data sources of geographical, venue-specific, social, linguistic aspects enables the site owners to detect users who are about to churn. After identifying those users who have a high probability of churning, one can employ gamification techniques such as badges and rewards to motivate them not to leave the service. For example, they could be awarded for additional reviews after a long period of inactivity.

10 CONCLUSION

In this paper, we studied the engagement patterns of producer-type users based on various aspects including geographical, linguistic, venue-specific, and social features. We performed a large-scale analysis of the representative LBSNs (*i.e.*, Yelp and Foursquare). We initially characterized user types on the employed large-scale datasets to focus our analysis on long-term producers who contribute the most UGC to the community among all user types. After that, we examined how long-term producers behave geographically in the offline real world and engage with the online community of LBSNs. First, in contrast to the human life course assumption, we found that users exhibit exploring behaviors until the end of their life in LBSNs. For example, they consistently travel to different locations at least 60% of all reviews for each life-stage. Second, we found that churning users and staying users show different patterns in four aspects. To name a few, staying users are more likely to travel to unexplored neighborhoods for reviewing and write reviews on diverse venues with more accumulated reviews. Besides, from the social aspect, we discovered that the churning of their friends profoundly influences long-term producers. Last but not least, we demonstrated the predictive models based on the insights derived from this work could successfully predict whether a long-term producer will leave the site. The classifiers learned with the proposed feature sets verified the effectiveness of those features.

There are many interesting directions that deserve further research. First of all, engagement patterns of newcomers using their location trajectories would be an important direction to study. We want to extract robust properties from user types, newcomers and long-term users, and develop advanced deep learning models to detect potential long-term users among newcomers. This research can help site maintainers to manage their user base from the influx of newcomers effectively. Second, analyzing user engagement using both users’ reviews and

check-in information can enhance our understanding of user behaviors in LBSNs. Finally, we want to incorporate more information such as demographic and personality traits of users so that we can identify a primary factor for each user type to churn. For that, we want to perform a comprehensive user survey with different types of users to investigate various motivations to stop contributing to the service.

ACKNOWLEDGMENTS

This research has been supported in part by project 16214817 from the Research Grants Council of Hong Kong and the 5GEAR project from the Academy of Finland.

REFERENCES

- [1] I. Adaji and J. Vassileva. 2015. Predicting Churn of Expert Respondents in Social Networks Using Data Mining Techniques: A Case Study of Stack Overflow. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*. 182–189. <https://doi.org/10.1109/ICMLA.2015.120>
- [2] Eytan Adar, Jaime Teevan, and Susan T. Dumais. 2008. Large Scale Analysis of Web Revisitation Patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1197–1206. <https://doi.org/10.1145/1357054.1357241>
- [3] Luca Maria Aiello and Nicola Barbieri. 2017. Evolution of Ego-networks in Social Media with Link Recommendations. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*. ACM, New York, NY, USA, 111–120. <https://doi.org/10.1145/3018661.3018733> event-place: Cambridge, United Kingdom.
- [4] Hadi Amiri and Hal Daume Iii. 2016. Short Text Representation for Detecting Churn in Microblogs. In *Thirtieth AAAI Conference on Artificial Intelligence*. <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12296>
- [5] Jaime Arguello, Brian S. Butler, Elisabeth Joyce, Robert Kraut, Kimberly S. Ling, Carolyn Rosé, and Xiaoqing Wang. 2006. Talk to Me: Foundations for Successful Individual-group Interactions in Online Communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 959–968. <https://doi.org/10.1145/1124772.1124916>
- [6] Valerio Arnaboldi, Marco Conti, Andrea Passarella, and Robin I. M. Dunbar. 2017. Online Social Networks and information diffusion: The role of ego networks. *Online Social Networks and Media* 1 (June 2017), 44–55. <https://doi.org/10.1016/j.osnem.2017.04.001>
- [7] Lars Backstrom, Dan Huttenlocher, Jon Kleinberg, and Xiangyang Lan. 2006. Group Formation in Large Social Networks: Membership, Growth, and Evolution. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '06)*. ACM, New York, NY, USA, 44–54. <https://doi.org/10.1145/1150402.1150412>
- [8] Jie Bao, Yu Zheng, and Mohamed F. Mokbel. 2012. Location-based and Preference-aware Recommendation Using Sparse Geo-social Networking Data. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '12)*. ACM, New York, NY, USA, 199–208. <https://doi.org/10.1145/2424321.2424348>
- [9] Jie Bao, Yu Zheng, David Wilkie, and Mohamed Mokbel. 2015. Recommendations in location-based social networks: a survey. *GeoInformatica* 19, 3 (July 2015), 525–565. <https://doi.org/10.1007/s10707-014-0220-8>
- [10] Gustavo E. A. P. A. Batista, Ronaldo C. Prati, and Maria Carolina Monard. 2004. A Study of the Behavior of Several Methods for Balancing Machine Learning Training Data. *SIGKDD Explor. Newsl.* 6, 1 (June 2004), 20–29. <https://doi.org/10.1145/1007730.1007735>
- [11] Hancheng Cao, Zhilong Chen, Fengli Xu, Yong Li, and Vassilis Kostakos. 2018. Revisitation in Urban Space vs. Online: A Comparison Across POIs, Websites, and Smartphone Apps. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4 (Dec. 2018), 156:1–156:24. <https://doi.org/10.1145/3287034>
- [12] Xinlei Chen, Yu Wang, Jiayou He, Shijia Pan, Yong Li, and Pei Zhang. 2019. CAP: Context-aware App Usage Prediction with Heterogeneous Graph Embedding. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 1 (March 2019), 4:1–4:25. <https://doi.org/10.1145/3314391>
- [13] Y. Chen, J. Hu, H. Zhao, Y. Xiao, and P. Hui. 2018. Measurement and Analysis of the Swarm Social Network With Tens of Millions of Nodes. *IEEE Access* 6 (2018), 4547–4559. <https://doi.org/10.1109/ACCESS.2018.2789915>
- [14] Yang Chen, Chenfan Zhuang, Qiang Cao, and Pan Hui. 2014. Understanding Cross-site Linking in Online Social Networks. In *Proceedings of the 8th Workshop on Social Network Mining and Analysis (SNAKDD'14)*. ACM, New York, NY, USA, 6:1–6:9. <https://doi.org/10.1145/2659480.2659498>
- [15] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and Mobility: User Movement in Location-based Social Networks (KDD '11). ACM, New York, NY, USA, 1082–1090. <https://doi.org/10.1145/2020408.2020579>
- [16] Dongho Choi, Chirag Shah, and Vivek Singh. 2016. Probing the Interconnections Between Geo-exploration and Information Exploration Behavior. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1170–1175. <https://doi.org/10.1145/2971648.2971694>
- [17] Cindy Chung and James Pennebaker. 2007. The Psychological Functions of Function Words. In *Social communication*. Psychology Press, New York, NY, US, 343–359.

- [18] Justin Cranshaw, Raz Schwartz, Jason I Hong, and Norman Sadeh. 2012. The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City. In *Sixth International AAAI Conference on Web and Social Media*. 8. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM12/paper/download/4682/4967>
- [19] Justin Cranshaw, Eran Toch, Jason Hong, Aniket Kittur, and Norman Sadeh. 2010. Bridging the Gap Between Physical Location and Online Social Networks. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*. ACM, New York, NY, USA, 119–128. <https://doi.org/10.1145/1864349.1864380>
- [20] Elizabeth M. Daly and Werner Geyer. 2011. Effective Event Discovery: Using Location and Social Information for Scoping Event Recommendations. In *Proceedings of the Fifth ACM Conference on Recommender Systems (RecSys '11)*. ACM, New York, NY, USA, 277–280. <https://doi.org/10.1145/2043932.2043982>
- [21] Cristian Danescu-Niculescu-Mizil, Robert West, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. No Country for Old Members: User Lifecycle and Linguistic Change in Online Communities. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13)*. ACM, New York, NY, USA, 307–318. <https://doi.org/10.1145/2488388.2488416>
- [22] Koustuv Dasgupta, Rahul Singh, Balaji Viswanathan, Dipanjan Chakraborty, Sougata Mukherjee, Amit A. Nanavati, and Anupam Joshi. 2008. Social Ties and Their Relevance to Churn in Mobile Telecom Networks. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology (EDBT '08)*. ACM, New York, NY, USA, 668–677. <https://doi.org/10.1145/1353343.1353424>
- [23] Gideon Dror, Dan Pelleg, Oleg Rokhlenko, and Idan Szepietor. 2012. Churn Prediction in New Users of Yahoo! Answers. In *Proceedings of the 21st International Conference on World Wide Web (WWW '12 Companion)*. ACM, New York, NY, USA, 829–834. <https://doi.org/10.1145/2187980.2188207>
- [24] Kritika D'Silva, Kasthuri Jayarajah, Anastasios Noulas, Cecilia Mascolo, and Archan Misra. 2018. The Role of Urban Mobility in Retail Business Survival. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3 (Sept. 2018), 100:1–100:22. <https://doi.org/10.1145/3264910>
- [25] Erik H. Erikson and Joan M. Erikson. 1998. *The Life Cycle Completed (Extended Version)*. W. W. Norton.
- [26] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. 249–256. <http://proceedings.mlr.press/v9/glorot10a.html>
- [27] Marta C. González, César A. Hidalgo, and Albert-László Barabási. 2008. Understanding individual human mobility patterns. *Nature* 453, 7196 (June 2008), 779–782. <https://doi.org/10.1038/nature06958>
- [28] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech Recognition with Deep Recurrent Neural Networks. *arXiv:1303.5778 [cs]* (March 2013). <http://arxiv.org/abs/1303.5778> arXiv: 1303.5778.
- [29] Reza Hadi Mogavi, Sujit Gujar, Xiaojuan Ma, and Pan Hui. 2019. HRCR: Hidden Markov-based Reinforcement to Reduce Churn in Question Answering Forums. In *Pacific Rim International Conference on Artificial Intelligence (PRICAI '19)*. Springer.
- [30] J. Hamari, J. Koivisto, and H. Sarsa. 2014. Does Gamification Work? – A Literature Review of Empirical Studies on Gamification. In *2014 47th Hawaii International Conference on System Sciences*. 3025–3034. <https://doi.org/10.1109/HICSS.2014.377>
- [31] William L. Hamilton, Justine Zhang, Cristian Danescu-Niculescu-Mizil, Dan Jurafsky, and Jure Leskovec. 2017. Loyalty in Online Communities. In *Eleventh International AAAI Conference on Web and Social Media*. <https://aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15710>
- [32] Gabriella M. Harari, Nicholas D. Lane, Rui Wang, Benjamin S. Crosier, Andrew T. Campbell, and Samuel D. Gosling. 2016. Using Smartphones to Collect Behavioral Data in Psychological Science: Opportunities, Practical Considerations, and Challenges. *Perspectives on psychological science : a journal of the Association for Psychological Science* 11, 6 (Nov. 2016), 838–854. <https://doi.org/10.1177/1745691616650285>
- [33] P. Hui and S. Buchegger. 2009. Groupthink and Peer Pressure: Social Influence in Online Social Network Groups. In *2009 International Conference on Advances in Social Network Analysis and Mining*. 53–59. <https://doi.org/10.1109/ASONAM.2009.17>
- [34] Hyunseok Hwang, Taesoo Jung, and Euiho Suh. 2004. An LTV model and customer segmentation based on customer value: a case study on the wireless telecommunication industry. *Expert Systems with Applications* 26, 2 (Feb. 2004), 181–188. [https://doi.org/10.1016/S0957-4174\(03\)00133-7](https://doi.org/10.1016/S0957-4174(03)00133-7)
- [35] Shan Jiang, Yingxiang Yang, Siddharth Gupta, Daniele Veneziano, Shounak Athavale, and Marta C. González. 2016. The TimeGeo modeling framework for urban mobility without travel surveys. *Proceedings of the National Academy of Sciences* 113, 37 (Sept. 2016), E5370–E5378. <https://doi.org/10.1073/pnas.1524261113>
- [36] Simon L. Jones, Denzil Ferreira, Simo Hosio, Jorge Goncalves, and Vassilis Kostakos. 2015. Revisitation Analysis of Smartphone App Use. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 1197–1208. <https://doi.org/10.1145/2750858.2807542>
- [37] Ewa Kaciewicz, James W. Pennebaker, Matthew Davis, Moonjee Jeon, and Arthur C. Graesser. 2014. Pronoun Use Reflects Standings in Social Hierarchies. *Journal of Language and Social Psychology* 33, 2 (March 2014), 125–143. <https://doi.org/10.1177/0261927X13502654>
- [38] Sanjay Ram Kairam, Dan J. Wang, and Jure Leskovec. 2012. The Life and Death of Online Groups: Predicting Group Growth and Longevity. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12)*. ACM, New York, NY, USA, 673–682. <https://doi.org/10.1145/2124295.2124374>

- [39] Marcel Karnstedt, Tara Hennessy, Jeffrey Chan, Partha Basuchowdhuri, Conor Hayes, and Thorsten Strufe. 2010. Churn in Social Networks. Springer, Boston, MA, 185–220. https://doi.org/10.1007/978-1-4419-7142-5_9
- [40] Riivo Kikas, Marlon Dumas, and Márton Karsai. 2013. Bursty egocentric network evolution in Skype. *Social Network Analysis and Mining* 3, 4 (Dec. 2013), 1393–1401. <https://doi.org/10.1007/s13278-013-0123-y>
- [41] Sunghwan Mac Kim, Kyo Kageura, James McHugh, Surya Nepal, Cécile Paris, Bella Robinson, Ross Sparks, and Stephen Wan. 2017. Twitter Content Eliciting User Engagement: A Case Study on Australian Organisations. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW '17 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 807–808. <https://doi.org/10.1145/3041021.3054237>
- [42] Gueorgi Kossinets and Duncan J. Watts. 2006. Empirical Analysis of an Evolving Social Network. *Science* 311, 5757 (Jan. 2006), 88–90. <https://doi.org/10.1126/science.1116869>
- [43] V. Kumar, Yashoda Bhagwat, and Xi (Alan) Zhang. 2015. Regaining “Lost” Customers: The Predictive Power of First-Lifetime Behavior, the Reason for Defection, and the Nature of the Win-Back Offer. *Journal of Marketing* 79, 4 (May 2015), 34–55. <https://doi.org/10.1509/jm.14.0107>
- [44] Young D. Kwon, Reza Hadi Mogavi, Ehsan Ul Haq, Youngjin Kwon, Xiaojuan Ma, and Pan Hui. 2019. Effects of Ego Networks and Communities on Self-Disclosure in an Online Social Network. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2019 (ASONAM '19)*. ACM, Vancouver, BC, Canada. <https://doi.org/10.1145/3341161.3342881>
- [45] William Labov. 1966. *The Social Stratification of English in New York City*. Center for Applied Linguistics.
- [46] Mounia Lalmas, Heather O'Brien, and Elad Yom-Tov. 2014. *Measuring User Engagement*. Morgan & Claypool Publishers.
- [47] Cliff Lampe, Rick Wash, Alcides Velasquez, and Elif Ozkaya. 2010. Motivations to Participate in Online Communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1927–1936. <https://doi.org/10.1145/1753326.1753616>
- [48] Quoc V. Le, Jiquan Ngiam, Adam Coates, Abhik Lahiri, Bobby Prochnow, and Andrew Y. Ng. 2011. On Optimization Methods for Deep Learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning (ICML '11)*. Omnipress, USA, 265–272. <http://dl.acm.org/citation.cfm?id=3104482.3104516> event-place: Bellevue, Washington, USA.
- [49] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (May 2015), 436–444. <https://doi.org/10.1038/nature14539>
- [50] Jure Leskovec, Lars Backstrom, Ravi Kumar, and Andrew Tomkins. 2008. Microscopic Evolution of Social Networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '08)*. ACM, New York, NY, USA, 462–470. <https://doi.org/10.1145/1401890.1401948>
- [51] Yuxuan Liang, Songyu Ke, Junbo Zhang, Xiuwen Yi, and Yu Zheng. 2018. GeoMAN: Multi-level Attention Networks for Geo-sensory Time Series Prediction. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden, 3428–3434. <https://doi.org/10.24963/ijcai.2018/476>
- [52] Shihan Lin, Rong Xie, Qinge Xie, Hao Zhao, and Yang Chen. 2017. Understanding User Activity Patterns of the Swarm App: A Data-driven Study (*UbiComp '17*). ACM, New York, NY, USA, 125–128. <https://doi.org/10.1145/3123024.3123086>
- [53] Zhiyuan Lin, Tim Althoff, and Jure Leskovec. 2018. I’ll Be Back: On the Multiple Lives of Users of a Mobile Activity Tracking Application. In *Proceedings of the 2018 World Wide Web Conference (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1501–1511. <https://doi.org/10.1145/3178876.3186062>
- [54] Yanchi Liu, Chuanren Liu, Xinjiang Lu, Mingfei Teng, Hengshu Zhu, and Hui Xiong. 2017. Point-of-Interest Demand Modeling with Human Mobility Patterns. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. ACM, New York, NY, USA, 947–955. <https://doi.org/10.1145/3097983.3098168>
- [55] Xinjiang Lu, Zhiwen Yu, He Du, Fei Yi, and Bin Guo. 2017. Discovery of Booming and Decaying Point-of-interest with Human Mobility Data. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers (UbiComp '17)*. ACM, New York, NY, USA, 137–140. <https://doi.org/10.1145/3123024.3123146>
- [56] Xinjiang Lu, Zhiwen Yu, Leilei Sun, Chuanren Liu, Hui Xiong, and Chu Guan. 2016. Characterizing the Life Cycle of Point of Interests Using Human Mobility Patterns. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1052–1063. <https://doi.org/10.1145/2971648.2971749>
- [57] Augusto Q. Macedo, Leandro B. Marinho, and Rodrygo L.T. Santos. 2015. Context-Aware Event Recommendation in Event-based Social Networks. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. ACM, New York, NY, USA, 123–130. <https://doi.org/10.1145/2792838.2800187>
- [58] Pablo Martí, Leticia Serrano-Estrada, and Almudena Nolasco-Cirugeda. 2019. Social Media data: Challenges, opportunities and limitations in urban studies. *Computers, Environment and Urban Systems* 74 (March 2019), 161–174. <https://doi.org/10.1016/j.compenvurbsys.2018.11.001>
- [59] Akhil Mathur, Nicholas D. Lane, and Fahim Kawsar. 2016. Engagement-aware Computing: Modelling User Engagement from Mobile Contexts. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM,

- New York, NY, USA, 622–633. <https://doi.org/10.1145/2971648.2971760>
- [60] Michael Mehaffy, Sergio Porta, Yodan Rofé, and Nikos Salingaros. 2010. Urban nuclei and the geometry of streets: The ‘emergent neighborhoods’ model. *URBAN DESIGN International* 15, 1 (March 2010), 22–46. <https://doi.org/10.1057/udi.2009.26>
- [61] Haim Mendelson and Ken Moon. 2018. Modeling Success and Engagement for the App Economy. In *Proceedings of the 2018 World Wide Web Conference (WWW ’18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 569–578. <https://doi.org/10.1145/3178876.3186123>
- [62] Ben Miroglia, David Zeber, Jofish Kaye, and Rebecca Weiss. 2018. The Effect of Ad Blocking on User Engagement with the Web. In *Proceedings of the 2018 World Wide Web Conference (WWW ’18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 813–821. <https://doi.org/10.1145/3178876.3186162>
- [63] M. C. Mozer, R. Wolniewicz, D. B. Grimes, E. Johnson, and H. Kaushansky. 2000. Predicting subscriber dissatisfaction and improving retention in the wireless telecommunications industry. *IEEE Transactions on Neural Networks* 11, 3 (May 2000), 690–696. <https://doi.org/10.1109/72.846740>
- [64] Dong Nguyen and Carolyn P. Rosé. 2011. Language Use As a Reflection of Socialization in Online Communities. In *Proceedings of the Workshop on Languages in Social Media (LSM ’11)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 76–85. <http://dl.acm.org/citation.cfm?id=2021109.2021119>
- [65] Guangli Nie, Wei Rowe, Lingling Zhang, Yingjie Tian, and Yong Shi. 2011. Credit card churn forecasting by logistic regression and decision tree. *Expert Systems with Applications* 38, 12 (Nov. 2011), 15273–15285. <https://doi.org/10.1016/j.eswa.2011.06.028>
- [66] Hartmut Obendorf, Harald Weinreich, Eelco Herder, and Matthias Mayer. 2007. Web Page Revisitation Revisited: Implications of a Long-term Click-stream Study of Browser Usage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’07)*. ACM, New York, NY, USA, 597–606. <https://doi.org/10.1145/1240624.1240719>
- [67] Richard J. Oentaryo, Ee-Peng Lim, David Lo, Feida Zhu, and Philips K. Prasetyo. 2012. Collective Churn Prediction in Social Network. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012) (ASONAM ’12)*. IEEE Computer Society, Washington, DC, USA, 210–214. <https://doi.org/10.1109/ASONAM.2012.44>
- [68] Jagat Sastry Pudipeddi, Leman Akoglu, and Hanghang Tong. 2014. User Churn in Focused Question Answering Sites: Characterizations and Prediction. In *Proceedings of the 23rd International Conference on World Wide Web (WWW ’14 Companion)*. ACM, New York, NY, USA, 469–474. <https://doi.org/10.1145/2567948.2576965>
- [69] Philipp Pushnyakov and Gleb Gusev. 2014. User Profiles Based on Revisitation Times. In *Proceedings of the 23rd International Conference on World Wide Web (WWW ’14 Companion)*. ACM, New York, NY, USA, 359–360. <https://doi.org/10.1145/2567948.2577380>
- [70] Al M. Rashid, Kimberly Ling, Regina D. Tassone, Paul Resnick, Robert Kraut, and John Riedl. 2006. Motivating Participation by Displaying the Value of Contribution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’06)*. ACM, New York, NY, USA, 955–958. <https://doi.org/10.1145/1124772.1124915>
- [71] Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. 2013. Early stopping and non-parametric regression: An optimal data-dependent stopping rule. *arXiv:1306.3574 [stat]* (June 2013). <http://arxiv.org/abs/1306.3574> arXiv: 1306.3574.
- [72] John C. Sherblom. 1990. Organization involvement expressed through pronoun use in computer mediated communication. *Communication Research Reports* 7, 1 (June 1990), 45–50. <https://doi.org/10.1080/08824099009359853>
- [73] Chaoming Song, Tal Koren, Pu Wang, and Albert-László Barabási. 2010. Modelling the scaling properties of human mobility. *Nature Physics* 6, 10 (Oct. 2010), 818–823. <https://doi.org/10.1038/nphys1760>
- [74] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of Predictability in Human Mobility. *Science* 327, 5968 (Feb. 2010), 1018–1021. <https://doi.org/10.1126/science.1177170>
- [75] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15 (2014), 1929–1958. <http://jmlr.org/papers/v15/srivastava14a.html>
- [76] Chenhao Tan and Lillian Lee. 2015. All Who Wander: On the Prevalence and Characteristics of Multi-community Engagement. In *Proceedings of the 24th International Conference on World Wide Web (WWW ’15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1056–1066. <https://doi.org/10.1145/2736277.2741661>
- [77] Zhen Tu, Yali Fan, Yong Li, Xiang Chen, Li Su, and Depeng Jin. 2019. From Fingerprint to Footprint: Cold-start Location Recommendation by Learning User Interest from App Data. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 1 (March 2019), 26:1–26:22. <https://doi.org/10.1145/3314413>
- [78] Marisa Affonso Vasconcelos, Saulo Ricci, Jussara Almeida, Fabrício Benevenuto, and Virgílio Almeida. 2012. Tips, Dones and Todos: Uncovering User Profiles in Foursquare (WSDM ’12). ACM, New York, NY, USA, 653–662. <https://doi.org/10.1145/2124295.2124372>
- [79] P. Wang, W. He, and J. Zhao. 2014. A Tale of Three Social Networks: User Activity Comparisons across Facebook, Twitter, and Foursquare. *IEEE Internet Computing* 18, 2 (March 2014), 10–15. <https://doi.org/10.1109/MIC.2013.128>
- [80] Y. Wang, Y. Guo, and Y. Chen. 2016. Accurate and early prediction of user lifespan in an online video-on-demand system. In *2016 IEEE 13th International Conference on Signal Processing (ICSP)*. 969–974. <https://doi.org/10.1109/ICSP.2016.7877974>

- [81] Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *Nature* 393, 6684 (June 1998), 440–442. <https://doi.org/10.1038/30918>
- [82] Canwen Xu, Jing Li, Xiangyang Luo, Jiaxin Pei, Chenliang Li, and Donghong Ji. 2019. DLocRL: A Deep Learning Pipeline for Fine-Grained Location Recognition and Linking in Tweets. In *The World Wide Web Conference (WWW ’19)*. ACM, New York, NY, USA, 3391–3397. <https://doi.org/10.1145/3308558.3313491> event-place: San Francisco, CA, USA.
- [83] Fengli Xu, Tong Xia, Hancheng Cao, Yong Li, Funing Sun, and Fanchao Meng. 2018. Detecting Popular Temporal Modes in Population-scale Unlabelled Trajectory Data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (March 2018), 46. <https://doi.org/10.1145/3191778>
- [84] Z. Xu, B. Chen, X. Meng, and L. Liu. 2017. Towards efficient detection of sybil attacks in location-based social networks. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. 1–7. <https://doi.org/10.1109/SSCI.2017.8285447>
- [85] Yuan Xuan, Yang Chen, Huiying Li, Pan Hui, and Lei Shi. 2016. LBSNShield: Malicious Account Detection in Location-Based Social Networks (CSCW ’16 Companion). ACM, New York, NY, USA, 437–440. <https://doi.org/10.1145/2818052.2869094>
- [86] Carl Yang, Xiaolin Shi, Luo Jie, and Jiawei Han. 2018. I Know You’ll Be Back: Interpretable New User Clustering and Churn Prediction on a Mobile Social Application. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD ’18)*. ACM, New York, NY, USA, 914–922. <https://doi.org/10.1145/3219819.3219821>
- [87] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudre-Mauroux. 2019. Revisiting User Mobility and Social Relationships in LBSNs: A Hypergraph Embedding Approach. In *The World Wide Web Conference on - WWW ’19*. ACM Press, San Francisco, CA, USA, 2147–2157. <https://doi.org/10.1145/3308558.3313635>
- [88] Dingqi Yang, Daqing Zhang, Zhiyong Yu, and Zhu Wang. 2013. A Sentiment-enhanced Personalized Location Recommendation System (HT ’13). ACM, New York, NY, USA, 119–128. <https://doi.org/10.1145/2481492.2481505>
- [89] Jiang Yang, Xiao Wei, Mark S. Ackerman, and Lada A. Adamic. 2010. Activity Lifespan: An Analysis of User Survival Patterns in Online Knowledge Sharing Communities. In *Fourth International AAAI Conference on Weblogs and Social Media*. <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/view/1466>
- [90] Yang Yang, Zongtao Liu, Chenhao Tan, Fei Wu, Yueting Zhuang, and Yafeng Li. 2018. To Stay or to Leave: Churn Prediction for Urban Migrants in the Initial Period. In *Proceedings of the 2018 World Wide Web Conference (WWW ’18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 967–976. <https://doi.org/10.1145/3178876.3186144>
- [91] Yiming Yang and Jan O. Pedersen. 1997. A Comparative Study on Feature Selection in Text Categorization. In *Proceedings of the Fourteenth International Conference on Machine Learning (ICML ’97)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 412–420. <http://dl.acm.org/citation.cfm?id=645526.657137>
- [92] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik-Lun Lee. 2011. Exploiting Geographical Influence for Collaborative Point-of-interest Recommendation. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR ’11)*. ACM, New York, NY, USA, 325–334. <https://doi.org/10.1145/2009916.2009962>
- [93] Jing Yuan, Yu Zheng, and Xing Xie. 2012. Discovering Regions of Different Functions in a City Using Human Mobility and POIs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’12)*. ACM, New York, NY, USA, 186–194. <https://doi.org/10.1145/2339530.2339561>
- [94] Amy X. Zhang, Anastasios Noulas, Salvatore Scellato, and Cecilia Mascolo. 2013. Hoodsquare: Modeling and Recommending Neighborhoods in Location-based Social Networks. *2013 International Conference on Social Computing* (Sept. 2013), 69–74. <https://doi.org/10.1109/SocialCom.2013.17> arXiv: 1308.3657.
- [95] Justine Zhang, William L. Hamilton, Cristian Danescu-Niculescu-Mizil, Dan Jurafsky, and Jure Leskovec. 2017. Community Identity and User Engagement in a Multi-Community Landscape. *Proceedings of the 11th International Conference On Web And Social Media, ICWSM 2017* 2017 (May 2017), 377–386. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5774974/>
- [96] Zengbin Zhang, Lin Zhou, Xiaohan Zhao, Gang Wang, Yu Su, Miriam Metzger, Haitao Zheng, and Ben Y. Zhao. 2013. On the Validity of Geosocial Mobility Traces. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*. ACM, New York, NY, USA, 11:1–11:7. <https://doi.org/10.1145/2535771.2535786>
- [97] Yu Zheng. 2011. Location-Based Social Networks: Users. In *Computing with Spatial Trajectories*. Springer, New York, NY, 243–276. https://doi.org/10.1007/978-1-4614-1629-6_8
- [98] Yin Zhu, Erheng Zhong, Sinno Jialin Pan, Xiao Wang, Minzhe Zhou, and Qiang Yang. 2013. Predicting User Activity Level in Social Networks. In *Proceedings of the 22Nd ACM International Conference on Information & Knowledge Management (CIKM ’13)*. ACM, New York, NY, USA, 159–168. <https://doi.org/10.1145/2505515.2505518>