

SAGE: Interactive State-aware Point-of-Interest Recommendation

Behrooz Omidvar-Tehrani, Sruthi Viswanathan, Frederic Roulland, Jean-Michel Renders
firstname.lastname@naverlabs.com
NAVER LABS Europe

ABSTRACT

Point-of-Interest (POI) recommendation is surfacing in many location-based services. User models are employed in these services to leverage historical check-ins and social links, and enable personalized and socialized POI recommendations. However these models often lack *interactivity* (incorporating user interactions) and *state-awareness*. This deficiency aggravates in cold start situations, where nearly no user information (historical check-ins and social graph) is available to generate effective recommendations. In this paper, we propose SAGE, an interactive state-aware POI recommendation system which tackles the aforementioned challenges by exploiting *look-alike groups* mined in public POI datasets, such as Foursquare and Yelp. SAGE reformulates the problem of POI recommendation as recommending explainable look-alike groups (and their POIs) which are in line with user's intent. SAGE frames the task of POI recommendation as an exploratory process where users interact with the system, and their interactions impact the way look-alike groups are picked out. Moreover, SAGE defines and employs *mindsets* which capture the actual state of the user and enforce the semantics of POI interestingness. Our experiments show that SAGE is an effective approach to capture interactivity and contextuality for recommending relevant look-alike groups and their POIs which are oriented towards the user's mindset.

ACM Reference Format:

Behrooz Omidvar-Tehrani, Sruthi Viswanathan, Frederic Roulland, Jean-Michel Renders. 2020. SAGE: Interactive State-aware Point-of-Interest Recommendation. In *WSDM Workshop SUM'20: State-based User Modelling, February 7, 2020, Houston, Texas, TX, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

There has been a meteoric rise in the use of location-based systems to benefit from services such as exploratory map browsing [1], localized advertising [2], and regional health-care [3]. Point-of-Interest (POI) recommendation is one of the most prominent applications of location-based systems which benefit both consumers and enterprises. It is shown in [4] that most users tend to visit POIs that they have not visited in the past 30 days. Hence the task of POI recommendation is to recommend users the POIs (e.g., restaurants,

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.

WSDMSUM '20, February 7, 2020, Houston, Texas, TX, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-9999-9/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

coffee shops, museums) that they may be interested in, but have never visited in a given time window.

While POI recommendation in general inherits the large body of work in the community of recommender systems, it also carries new constraints and challenges that might not be the case for a traditional recommender, such as spatial distance semantics between POIs and user interactions on maps. We believe that an ideal POI recommendation approach should capture the following aspects.

A1: Personalization. First, POI recommendations should be personalized, i.e., the results should be based on user preferences captured in form of user's historical check-ins and interests ([5, 6], to name a few).

A2: Socialization. People trust look-alike users and base their decisions on what people like them have appreciated before [7]. Hence the POI recommendation should also incorporate social aspects and reflect the preferences of others similar to the user. Socialization has been addressed in the literature ([8, 9], to name a few), where information encapsulated in location-based social networks (LBSN) are employed to predict user's preferences using link-based methods [10].

A3: Interactivity. Beyond being personalized, the POI recommendation system should also be *exploratory* to incorporate user's interactions with the system and customize recommended POIs accordingly [11].

A4: State-awareness. The POI recommendation should also capture the current state of the user. While the literature focuses mostly on contextual features of the user state (such as time and location), actual situation and intents of users have received less attention.

To the best of our knowledge, no POI recommendation approach in the literature addresses all the aforementioned aspects simultaneously, due to the following challenges.

C1: Cold start and data sparsity. The problem of cold start arises when a user with a limited history of check-ins asks for recommendations. Also data sparsity refers to the lack of data for identifying similarities between users. Many users employ POI services such as Foursquare and Yelp without signing in. As a result, no social graph (i.e., friendship relations) can be retrieved. A typical recommendation system which relies on historical check-ins and user similarities for personalized and socialized recommendations (**A1** and **A2**, respectively) is unable to output results in the presence of cold start and data sparsity. There are two kinds of users who may cause a cold start: a new user with no history, and a user with privacy concerns who does not want his/her data to be exploited (as commonly done in mobile cloud services [12]). As both these cases are realistic, the recommendation paradigm should be redesigned to incorporate them.

C2: Interpretation of interactions. Most POI recommendation systems assume the process to be *one-shot*, where the user enters the system with a clear unambiguous intent, and the system returns the most interesting POIs related to that intent. In practice, this architecture is not realistic anymore. Users need to be inside the loop and interact with the system to gradually build their intent. The challenge with multi-shot recommendation systems (**A3**) is that it is not clear how user interactions with the system should influence the user state and the recommendation strategy.

C3: State integration. It is challenging to integrate the user state into the recommendation process (**A4**). The state is not limited to contextual features such as time and location, but also situational features such as the user's mindset at the time of receiving recommendations. For instance, different POIs should be recommended in case the user is hungry, or in case he/she is seeking some personal relaxation time (i.e., me time).

C4: Explainability. Users may not trust in what they get from the recommender due to the cold start problem (**A1**) and interactions with the system (**A3**). Hence it is of critical importance to let users know *why they receive certain POIs as recommendation results*. The challenge of recommendation transparency is a long standing problem, and some approaches employ textual features, such as aspect-based sentiment analysis to explain recommendations [13].

To collectively address the challenges **C1** to **C4**, we propose SAGE, an interactive state-aware POI recommendation system based on look-alike groups. The intuition behind SAGE is as follows: while it is assumed that no data is available on the user (i.e., cold start), POI recommendations can be obtained by finding look-alike groups in publicly available POI datasets such as Foursquare and Gowalla. It is shown in the literature that users trust peers and get inspired by them for decision making [7]. The recommended POIs are explainable using their associated groups, e.g., “the group of photoholics tends to visit La Butte in the 18th arrondissement of Paris”, and “the group of food lovers tends to visit the restaurant les Apotres de Pigalles, in the same region.” The user will then interact with those groups to detect with which group he/she identifies. As a result of this interaction, new groups will be mined, to align with the user’s intent. This iterative process ensures that groups and their POIs reflect user’s preferences. It is to note that SAGE discovers user’s intents and aligns recommendations accordingly, *without the need of any historical check-in data from the user*. The following example describes how SAGE is employed in practice.

EXAMPLE 1. Lindsey is visiting Paris as a tourist. She is walking in the area of the Pompidou center. After 30 minutes of walking, she gets tired and asks SAGE for “me time” recommendations, to find POIs in her vicinity (the dashed circle in Figure 1) within which she can sit and relax. The “me time” option is a signal of her situational state. Lindsey is concerned about privacy and does not share any historical check-in data with the system (hence causing cold start). SAGE outputs three user groups¹ related to Lindsey’s intent (i.e., having me-time), and top-three POIs for each group (Figure 1). She looks at group descriptions to see where she finds some doppelgängers among the group members. Being a social person, Lindsey shows interest in the yellow group, i.e., visitors who have many friends (i.e., social visitors like her), and tend

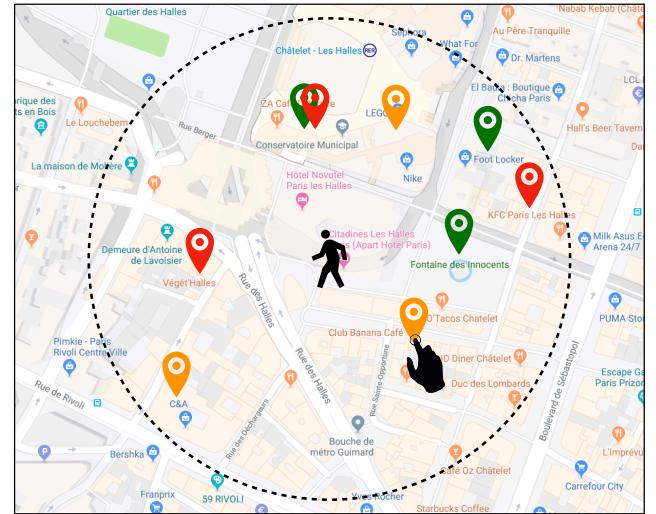


Figure 1: SAGE in practice. Green POIs are associated to a group of visitors who tend to visit historical landmarks. Red POIs relate to visitors who tend to visit Asian Food restaurants. Also yellow POIs relate to visitors who have many friends and tend to visit coffee shops and American restaurants.

to visit coffee shops and American restaurants. This motivates her to refine her intent and seek a relaxing place where she can also eat in. She interacts with the system and asks where people usually eat in the neighborhood. Hence SAGE returns another three groups to satisfy Lindsey’s intent. This helps her to make up her mind and go to an American Burger restaurant.

The above example shows that users can get actionable POI recommendations by interacting with look-alike groups. It also shows how an exploratory process tracks user states and helps users to refine their ambiguous needs and finalize their decisions. The way users can specify their mindsets (e.g., “me time” and “I’m hungry” in the example) enables users enforce their state to the recommendation system (i.e., tackling **C3**) and bias the results towards what they are really interested to receive.

In summary, we propose the following contributions towards an interactive state-aware POI recommendation system called SAGE.

- We address the cold start and explainability challenges by employing look-alike user data as a proxy for user preferences. We reformulate the problem of POI recommendation, as recommending explainable groups (and their POIs) which are in line with user’s intent.
- We consider “recommendation” as an exploratory process where the user interacts with the system, and his/her interactions impact the way look-alike groups are selected out.
- We introduce and formalize the notion of “mindsets”, which captures situational aspects of the user state such as actual situation and intent. We mention how interestingness of POIs is maximized using mindsets.

¹We will explain in Section 4 how groups are mined.

Utility function	Description
$popularity(P)$	normalized average number of check-ins of P .
$prestige(P)$	normalized average rating score of P .
$recency(P)$	inverse difference between the current date and the average insertion date of P .
$coverage(P)$	the area of a polygon induced by the geographical location of POIs in P normalized by the area of the city.
$surprisingness(P)$	normalized Jaccard distance between POI categories of P and POI categories of the bookmarked POIs by the user \mathcal{P}_μ .
$diversity(P)$	normalized Jaccard distance between sets of POI categories in P .
$size(P)$	normalized average radius of POIs in P .

Table 1: POI utility functions ($P \subseteq \mathcal{P}$).

- In an extensive set of quantitative and qualitative experiments, we show the efficiency and effectiveness of SAGE.

Outline. In Section 2, we provide the data model for interactive state-aware POI recommendation. In Section 3, we formally define our problem. In Section 4, we describe SAGE architecture and its underlying algorithm. Section 5 presents detailed experiments. The related work is provided in Section 6. Last, we conclude and discuss the future directions in Section 7.

2 DATA MODEL

We consider a user μ asking for POI recommendations. We also denote μ 's portfolio as a set \mathcal{P}_μ of POIs that μ is interested in. Following our assumption of cold start (**C1**), we consider that initially $\mathcal{P}_\mu = \emptyset$. Additionally, we consider a POI dataset $\mathcal{D} = \langle \mathcal{U}, \mathcal{P} \rangle$ with a set of visitors \mathcal{U} and a set of POIs \mathcal{P} .

Visitors. For a visitor $u \in \mathcal{U}$, the set $u.\text{demogs}$ contains tuples of the form $\langle d, v \rangle$ where d is a demographic attribute (e.g., age, gender, number of trips, number of check-ins), and $v \in \text{domain}(d)$. Also the set $u.\text{checkins}$ contains tuples of the form $\langle p, t \rangle$ which represents that u has visited a POI $p \in \mathcal{P}$ at time t .

POIs. A POI $p \in \mathcal{P}$ is defined as a tuple $p = \langle loc, att \rangle$ where $p.loc$ is itself a tuple $\langle lat, lon \rangle$ (latitude and longitude, respectively) which defines where p is situated geographically. The set $p.att$ is a set of tuples of the form $\langle a, v \rangle$ which denotes that the POI has the value v for the attribute a , such that $v \in \text{domain}(a)$.

We measure the interestingness of POIs using *POI utility* functions. A more interesting POI has higher chances to be recommended to the user μ . A POI utility function $f : 2^{\mathcal{P}} \mapsto [0, 1]$ returns a value between 0 and 1 which reflects the extent of interestingness for one or several POIs [14]. Table 1 lists POI utility functions that we employ in this work. We define all utility functions as maximization objectives.

Mindset label	Description
$m_1: \text{I'm new here}$	towards touristic POIs about the popular attractions in the city.
$m_2: \text{surprise me}$	towards POIs which haven't been visited before by the user and are uncommon (seldom visited)
$m_3: \text{let's workout}$	towards POIs related to physical exercises like swimming pools, parks, gyms, and mountains
$m_4: \text{me time}$	towards POIs related to activities to treat oneself and be pursued solo to unwind and relax
$m_5: \text{I'm hungry}$	towards getting faster access to food-related POIs nearby
$m_6: \text{let's learn}$	towards POIs such as museums, libraries and cultural landmarks
$m_7: \text{hidden gems}$	towards small intriguing local POIs that are highly rated but not necessarily popular

Table 2: Mindsets

3 PROBLEM DEFINITION

The problem of interactive state-aware POI recommendation builds on two core assumptions. First, we assume that the user μ is in an *exploratory setting* and does not necessarily have a clear idea of his/her needs [15], and he/she is going to sharpen his/her intent in several iterations. Second, we conjecture that *look-alike user data* is a good proxy to gain user preferences [16], in the absence of historical check-in data.

Exploratory setting. A core concept in an exploratory POI recommendation setting is “user state”, which is often materialized using contextual features such as current time and location of the user, $c_\mu = \langle loc, time \rangle$. An additional dimension of the user state is “mindsets”, i.e., actual situation and intents of the user. Mindsets should reflect the way interestingness of POIs are computed based on user’s intent. While online services such as AROUNDME² enable users to explore their nearby region by selecting explicit POI categories (e.g., museums), mindsets capture the intents of users (e.g., “let’s learn”) which are more challenging to capture. Table 2 lists the set of mindsets that we consider in this work.

A mindset m is a tuple $m = \langle label, func() \rangle$, where $label$ provides a short description of the mindset, and $func()$ defines semantics of POI interestingness. For instance, in case $m.label = "I'm hungry"$, $m.func()$ is formulated in a way to increase the interestingness score of restaurants and coffee shops. Also in case $m.label = "let's learn"$, $m.func()$ biases museums, libraries and cultural landmarks. Given a mindset m , the function $m.func()$ is defined as follows.

$$m.func(P, \mu) = \frac{\sum_{f_i(P) \in \mathcal{F}} b_{i,m} w_{i,\mu} f_i(P)}{\sum_{f_i(P) \in \mathcal{F}} b_{i,m} w_{i,\mu}} \quad (1)$$

In Equation 1, $f_i(P)$ is a utility function (see Table 1), and $b_{i,m}$ and $w_{i,\mu}$ are the prior and user-specific weight of $f_i(P)$ for the

²<http://www.aroundmeapp.com>

	popularity	prestige	recency	coverage	surprisingness	diversity	size
<i>m</i> ₁ : I'm new here	0.25	0.25	0.10	0.15	0.00	0.25	0.00
<i>m</i> ₂ : surprise me	0.25	0.20	0.00	0.00	0.25	0.20	0.10
<i>m</i> ₃ : let's workout	0.35	0.40	0.00	0.25	0.00	0.00	0.00
<i>m</i> ₄ : me time	0.20	0.20	0.00	0.20	0.00	0.20	0.20
<i>m</i> ₅ : I'm hungry	0.00	0.40	0.20	0.40	0.00	0.00	0.00
<i>m</i> ₆ : let's learn	0.30	0.30	0.00	0.20	0.00	0.20	0.00
<i>m</i> ₇ : hidden gems	0.30	0.30	0.15	0.00	0.00	0.00	0.25

Table 3: Priors in mindsets. Non-zero values are highlighted.

mindset m , respectively. Hence the mindset function is defined as a normalized weighted sum over the priors and weights. Priors reflect the importance of a utility function for a mindset. In case $b_{i,m} = 0$, it means that f_i has no influence on the mindset m . On the contrary, in case $b_{i,m} = 1$, it means that the mindset m is defined solely based on f_i . The weights, on the other hand, are user-centric parameters and reflect the importance of a utility function for the user. A user may have more interest in popularity than coverage. The weights are assumed to shape up when the user interacts with the system. Given the set of all possible user-specific weights W , we initially set $\forall w \in W, w = 1.0$. In case $\mathcal{P}_\mu \neq \emptyset$, we set $w_{i,\mu} = f_i(\mathcal{P}_\mu)$. While the weights are dynamic and changes per user, priors can be learned offline and stay unchanged at the online execution. Table 3 shows priors in mindsets, which we acquired from an in-depth qualitative UX research study to understand cognitive needs of users in a POI recommendation system.

Look-alike user data. There exist various publicly available POI datasets, such as Yelp, TripAdvisor, Foursquare, and Gowalla, structured as $\mathcal{D} = \langle \mathcal{U}, \mathcal{P} \rangle$. To build look-alike relations in the POI dataset, we build “visitor groups” which aggregate a set of visitors with common demographics and/or POIs [16]. Visitor groups are obviously virtual and group members do not necessarily know each other. In other words, members of a group are “location friends” (who have checked in the same places) and not necessarily “social friends” (who are socially connected in an LBSN) [9]. A visitor group is a triple $g = \langle \text{members}, \text{demogs}, \text{POIs} \rangle$ where $g.\text{members} \subseteq \mathcal{U}$, $\forall u \in g.\text{members}, \forall \langle a, v \rangle \in g.\text{demogs}, \langle a, v \rangle \in u.\text{demogs}$, and $\forall u \in g.\text{members}, \forall p \in g.\text{POIs}, \exists \langle p, t \rangle \in u.\text{checkins}$.

Problem definition. We define the problem of interactive state-aware POI recommendation as follows. Given a user μ and his/her affiliated context $c_\mu = \langle \text{loc}, \text{time} \rangle$, a mindset $m = \langle \text{label}, \text{func}() \rangle$, a radius r , and integers k and k' , the problem is to find top- k groups G and k' POIs for each group in G , such that the following conditions are met.

- (i) $\forall g \in G, \mathcal{P}_\mu \cap g.\text{POIs} \neq \emptyset \oplus \mathcal{P}_\mu = \emptyset$;
- (ii) $\forall g \in G, \forall p \in g.\text{POIs}, \text{distance}(p, c_\mu.\text{loc}) \leq r$;
- (iii) $\sum_{g \in G} m.\text{func}(g.\text{POIs}, \mu)$ is maximized.

The first two conditions ensure that groups are relevant to the user and in vicinity of the user’s location. Considering the cold start assumption, we neutralize condition (i) in case $\mathcal{P}_\mu = \emptyset$. The third condition applies the input mindset to groups, and verifies whether POIs are maximally in line with the mindset.

4 OUR APPROACH

SAGE is a session-based system which begins with an ambiguous user’s intent for POI recommendation, and ends when he/she is satisfied with the resulting POIs. Each session consists of a finite sequence of iterations which captures interactions with the user. A new iteration begins by defining a mindset (which may remain the same as the previous iteration), which then results in k relevant groups and k' POIs for each group. At the end of each iteration, the user is free to bookmark some of the recommended POIs to be added to \mathcal{P}_μ . Hence there are two types of feedback that the user can provide to the system: *the mindset* (which may stay unchanged between consecutive iterations) and *POI bookmarks*. This multi-shot architecture contradicts most traditional single-shot POI recommendation approaches, by incorporating *user interactions* in the recommendation (i.e., addressing C2).

Algorithm 1 describes the flow of each iteration in SAGE. First, the system finds all nearby POIs which are at most r kilometers/miles far from the user (line 1). Given the set of nearby POIs $P \subseteq \mathcal{P}$, SAGE then retrieves all check-ins whose POI is in P (line 2). Then the system mines groups among checked-in visitors denoted as G^* (line 3). Given that $|G^*| \gg k$, the system finds k groups $G \subset G^*$ (s.t., $|G| = k$) which collectively maximize the mindset function (line 4). Finally, SAGE picks top- k' POIs for each group which are visited by the majority of the group members (line 5). Note that although the contextual and situational features of the user state are provided as input to Algorithm 1, it still works in cold start settings, as no historical check-ins or social graph of μ is used by the algorithm. In addition, one direction of our future work is to build a classifier to automatically predict mindsets.

Mining look-alike groups. Explainability is a crucial need in cold start settings. To address the challenge of explainability (C4), we aim to find *describable groups* which identify a set of visitors checking in a set of POIs. For this aim, we employ Frequent Itemset Mining (FIM) technique, where each group is a frequent itemset, and items are common demographic attributes and POIs of group members. While groups can be discovered in myriad ways [16], we choose FIM to obtain describable groups with overlaps, so that visitors can be a member of more than one group and be described in different ways. We employ the Apriori algorithm [17] to mine groups. The algorithm is known to be inefficient for large number of items. In SAGE, we are able to perform the mining process on-the-fly, thanks to the neighborhood filters preceding the algorithm. In other words, the algorithm mines groups only for visitors with

Algorithm 1: SAGE Algorithm

Input: Visitors \mathcal{U} and POIs \mathcal{P} , user context $c_\mu = \langle \text{loc}, \text{time} \rangle$, radius r , mindset m , number of groups k , number of POIs per group k'

Output: Groups G and their POIs P_G

- 1 $P \leftarrow \text{nearby_POIs}(\mathcal{P}, c_\mu.\text{loc}, r)$
- 2 $H \leftarrow \text{checkins_of}(P, c_\mu)$
- 3 $G^* \leftarrow \text{mine_groups}(\mathcal{U}, H)$
- 4 $G \leftarrow \text{maximize}(\mu, G^*, k, m)$
- 5 **for** each group $g \in G$ **do** $P_G.append(\text{top_POIs}(g, k'))$
- 6 **return** G, P_G

check-ins in the vicinity of the user. Hence the size of the visitor set is drastically reduced compared to $|\mathcal{U}|$.

Maximizing mindsets. Not all groups are equally interesting to the user. We need to pick k groups out of all mined groups which are in line with the mindset requested by the user. This will tackle the challenge of context integration (C3). Each mindset is associated to a function which is a set of utility functions combined in a linear fashion with the priors and user-specific weights (Equation 1). The mindset function admits as input a set of POIs, and returns a value in the range $[0, 1]$. Given a mindset m and a group g , we measure the utility of g regarding m 's functionality as $\text{group_utility}(g) = m.\text{func}(g.\text{POIs})$. Given the space of all group utility values, the problem is to find k groups with the largest values of group utility. As each mindset function is constructed as a combination of several utility functions, maximizing mindset functions is a multi-objective optimization problem in nature. However, we employ a simple scalarization approach in SAGE using the priors and weights to reduce the complexity of the problem to single-objective optimization. In SAGE, we employ the greedy-style optimization algorithm in [18] to maximize mindset functions.

5 EXPERIMENTS

In this work, we use Gowalla dataset, collected from a popular LBSN with 36,001,959 check-ins of 319,063 visitors over 2,844,076 POIs. GOWALLA is among the few datasets that provide attributes both for visitors and POIs. Hence we can form groups containing both demographic attributes and POIs. This increases the explainability of groups, and enables users find out which group they identify in. In this section, we evaluate the overall algorithmic behavior of SAGE by simulating interaction sessions and reporting the Hit Ratio measure. We also evaluate the user-centered aspects of SAGE using an extensive user study.

Simulation study. Our goal is to examine the overall behavior of SAGE in tackling the cold start problem, and providing interactivity, state-awareness, and explainability in POI recommendation. To remove the influence of human decisions from the exploratory process, we simulate interactions (by picking a random user μ in Gowalla, and a random check-in from $\mu.\text{checkins}$ as μ 's actual state) and report the Hit Ratio $HR@N$ for each simulated session: $HR@N = \sum_{i=1}^S (\sum_{j=1}^N \mathbb{1}(i, j, \mu)) / S$, where N is the number of iterations, S is the number of sessions (we set it to 100), and $\mathbb{1}(i, j, \mu)$

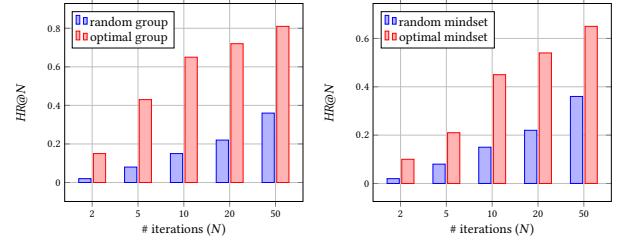


Figure 2: HR values for SAGE simulation.

is a function which returns “1” if there is at least one POI in common with $\mu.\text{checkins}$ in the iterations j of the session i . For a more realistic simulation, we filter out μ 's check-ins with more than 48 hours of time difference with μ 's state. Note that among several evaluation measures such as NDCG and MRR, we choose to report Hit Ratio because SAGE's recommendations are set-based and do not necessarily enforce rank semantics.

Figure 2 illustrates HR values using different strategies of group selection (left) and mindset selection (right). At each iteration, a group g is either picked at random (i.e., purely exploratory), or by maximizing $\text{Cosine}(g.\text{dem-ogs}, \mu.\text{demogs})$. Also a mindset is either picked at random, or by maximizing $m.\text{func}()$ for the POIs of the selected group g in the previous iteration (i.e., picking the most interesting mindset). Regarding the strategy of group selection, we observe that optimal groups increase HR by 38.8% on average. HR grows to values larger than 50% after only 10 iterations, and it reaches to 82% at 50 iterations. This confirms our assumption that look-alike groups function as a good proxy to gain user preferences. Regarding the mindset selection strategy, we observe that optimal mindsets increase HR by 22.4%. HR grows to values close to 50% after 10 iterations, and it reaches to 65% at 50 iterations.

User study. In this part of the experiments, we discuss human-oriented aspects of SAGE. We perform an extensive between-subject user study in Amazon Mechanical Turk³ to measure the effectiveness of employing look-alike groups and mindsets in SAGE. We recruited 753 participants in AMT and forwarded them to a Survey Monkey⁴ questionnaire to answer different questions about the functionality of our proposed system. In the first part of our study, we compare SAGE with a baseline, i.e., GOOGLE MAP EXPLORE, an interactive POI exploration and recommendation system. For a given region, the participant observes the results of SAGE and the baseline side-by-side, and should decide which set of POI recommendations he/she finds more useful. We found out that 59% of the participants prefer SAGE over its competitor. The a-priori familiarity of the participants with Google services was an influencing factor in this study, which led many participants to vote for their “comfort zone” technological method.

Moreover, we perform an independent study by describing an intuitive contextual state (e.g., “an evening in Paris”) and a mindset to the participant, and asking his/her opinion about the usefulness of SAGE's output. We pose two viewpoints, $v1$ and $v2$, and the participant expresses his/her agreement with each viewpoint in a

³<https://www.mturk.com>

⁴<https://www.surveymonkey.com>

Likert scale from 1 (totally disagreed) to 5 (totally agreed): ($v1$) the group-based results are relevant to the selected mindset, ($v2$) the groups help the participant understand why he/she receives the POIs as recommendation results. For both viewpoints, we observe that the supremacy of the agreement vote is statistically significant (72.97% for $v1$ and 72.65% for $v2$). This validates our hypotheses on the relevance of group-based recommendations to the mindset (i.e., viewpoint $v1$) and the understandability and explainability of recommended POIs using groups (i.e., viewpoint $v2$). We conducted one-way repeated ANOVA for the both results, and obtained the F statistics of 311.97 and 731.30 with the significance level of 0.05, for $v1$ and $v2$, respectively.

6 RELATED WORK

Traditionally, the problem of POI recommendation is defined as learning users' implicit preferences according to user's historical check-ins [19]. To solve this problem, most approaches in the literature employ memory-based and model-based Collaborative Filtering (CF) as the *de facto* approach, where the check-ins matrix is used for learning preferences.

Mindsets. Category-based search interfaces [20] capture explicit needs of users in the form of categories (e.g., selecting POIs of the category "historical landmarks"). Realistic scenarios often contain ambiguous needs and intents, where users seek to disambiguate in an iterative process. The only possible iteration in traditional search paradigms is to restart a search with another category. Exploratory travel interfaces [21] have been found to enhance user experience in POI exploration with serendipity measures. The transitory search interface [15] helps users discover activities in the city, empowering them to arrive at insightful results using slider continuums. In SAGE, we employ "mindsets", which is an intuitive way of capturing user's implicit intent as a situational state. Mindsets reflect the needs of users, and align to users' preferences during iterations.

Look-alike groups. To tackle the challenge of cold start, many approaches are proposed to enrich a sparse check-in matrix with social aspects, such as friendship links [22]. The assumption is that users may be more interested to visit places that their friends visited in the past [23]. In [22], a friendship-based CF is proposed. In [24], the popular Matrix Factorization model is extended with social regularization. Moreover, friendship links are exploited in [25] to build friend groups using Community Detection techniques. In SAGE, we extend the domain of cold start to social aspects, i.e., no friendship link is available for users. SAGE builds explainable look-alike groups without relying on any social aspects or check-ins of the user. Our hypothesis is that users are interested to visit places which people similar to them have visited before, which we validated in form of a user study.

7 CONCLUSION

We present SAGE, an interactive state-aware POI recommendation system based on look-alike groups, which tackles the common challenges of cold start, interactivity, state-awareness, and explainability. We introduce the notion of "mindsets" which extends the scope of user state, and captures actual situation and intents of the user. In an extensive set of experiments, we show that SAGE achieves a Hit Ratio higher than 50% only after 10 iterations. We

also showed the effectiveness of look-alike groups and mindsets for POI recommendation in an extensive user study.

REFERENCES

- [1] Leilani Battle, Remco Chang, and Michael Stonebraker. Dynamic prefetching of data tiles for interactive visualization. In *Proceedings of the 2016 International Conference on Management of Data*, pages 1363–1375. ACM, 2016.
- [2] Kaiyu Feng, Gao Cong, Sourav S Bhowmick, Wen-Chih Peng, and Chunyan Miao. Towards best region search for data exploration. In *CIKM*. ACM, 2016.
- [3] Cicer A. L. Pahin, Behrooz Omidvar-Tehrani, Sihem Amer-Yahia, Valerie Siroux, Jean-Louis Pepin, Jean-Christian Borel, and Joao Comba. COVIZ: A system for visual formation and exploration of patient cohorts. In *VLDB*, 2019.
- [4] Anastasios Noulas, Salvatore Scellato, Neal Lathia, and Cecilia Mascolo. A random walk around the city: New venue recommendation in location-based social networks. In *SOCIALCOM-PASSAT*. IEEE, 2012.
- [5] Bin Liu, Yanjie Fu, Zijun Yao, and Hui Xiong. Learning geographical preferences for point-of-interest recommendation. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1043–1051. ACM, 2013.
- [6] Hongzhi Yin, Bin Cui, Yizhou Sun, Zhiting Hu, and Ling Chen. Lcars: A spatial item recommender system. *ACM Transactions on Information Systems (TOIS)*, 32(3):11, 2014.
- [7] Fan Du, Catherine Plaisant, Neil Spring, and Ben Shneiderman. Finding similar people to guide life choices: Challenge, design, and evaluation. In *CHI*. ACM, 2017.
- [8] Xutao Li, Gao Cong, Xiao-Li Li, Tuan-Anh Nguyen Pham, and Shonali Krishnaswamy. Rank-geofm: A ranking based geographical factorization method for point of interest recommendation. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 433–442. ACM, 2015.
- [9] Huayu Li, Yong Ge, Richang Hong, and Hengshu Zhu. Point-of-interest recommendations: Learning potential check-ins from friends. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 975–984. ACM, 2016.
- [10] Hao Wang, Manolis Terrovitis, and Nikos Mamoulis. Location recommendation in location-based social networks using user check-in data. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 374–383. ACM, 2013.
- [11] Rishabh Mehrotra, Mounia Lalmas, Doug Kenney, Thomas Lim-Meng, and Golli Hashemian. Jointly leveraging intent and interaction signals to predict user satisfaction with slate recommendations. In *The World Wide Web Conference*, pages 1256–1267. ACM, 2019.
- [12] Hui Suo, Zuoohua Liu, Jiafu Wan, and Keliang Zhou. Security and privacy in mobile cloud computing. In *2013 9th International Wireless Communications and Mobile Computing Conference (IWCMC)*, pages 655–659. IEEE, 2013.
- [13] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 241–250. ACM, 2000.
- [14] Liqiang Geng and Howard J Hamilton. Interestingness measures for data mining: A survey. *ACM Computing Surveys (CSUR)*, 38(3):9, 2006.
- [15] Jeni Paay, Jesper Kjeldskov, Mikael B Skov, Per M Nielsen, and Jon Pearce. Discovering activities in your city using transitory search. In *MobileHCI*. ACM, 2016.
- [16] Behrooz Omidvar-Tehrani and Sihem Amer-Yahia. User group analytics: Survey and research opportunities. *TKDE*, 2019.
- [17] Rakesh Agrawal, Ramakrishnan Srikant, et al. Fast algorithms for mining association rules. In *Proc. 20th int. conf. very large data bases, VLDB*, volume 1215, pages 487–499, 1994.
- [18] Behrooz Omidvar-Tehrani, Sihem Amer-Yahia, and Ria Mae Borromeo. User group analytics: hypothesis generation and exploratory analysis of user data. *VLDB J.*, 28(2):243–266, 2019.
- [19] Jie Bao, Yu Zheng, David Wilkie, and Mohamed Mokbel. Recommendations in location-based social networks: a survey. *GeoInformatica*, 2015.
- [20] Marti Hearst. *Search user interfaces*. Cambridge university press, 2009.
- [21] Sruthi Viswanathan, Cecile Boulaire, and Antonietta Maria Grasso. Ageing clouds: Novel yet natural support for urban exploration. In *DIS*. ACM, 2019.
- [22] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik-Lun Lee. Exploiting geographical influence for collaborative point-of-interest recommendation. In *SIGIR*. ACM, 2011.
- [23] Diana Mok, Barry Wellman, and Juan Carrasco. Does distance matter in the age of the internet? *Urban Studies*, 47(13):2747–2783, 2010.
- [24] Mohsen Jamali and Martin Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *RecSys*. ACM, 2010.
- [25] Deepika Lalwani, Durvasula VLN Somayajulu, and P Radha Krishna. A community driven social recommendation system. In *Big Data*. IEEE, 2015.