

Content-Aware Point of Interest Recommendation on Location-Based Social Networks

Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu

Computer Science and Engineering
Arizona State University
{Huiji.Gao, Jiliang.Tang, Xia.Hu, Huan.Liu}@asu.edu

Abstract

The rapid urban expansion has greatly extended the physical boundary of users' living area and developed a large number of POIs (points of interest). POI recommendation is a task that facilitates users' urban exploration and helps them filter uninteresting POIs for decision making. While existing work of POI recommendation on location-based social networks (LBSNs) discovers the spatial, temporal, and social patterns of user check-in behavior, the use of content information has not been systematically studied. The various types of content information available on LBSNs could be related to different aspects of a user's check-in action, providing a unique opportunity for POI recommendation. In this work, we study the content information on LBSNs w.r.t. POI properties, user interests, and sentiment indications. We model the three types of information under a unified POI recommendation framework with the consideration of their relationship to check-in actions. The experimental results exhibit the significance of content information in explaining user behavior, and demonstrate its power to improve POI recommendation performance on LBSNs.

Introduction

The rapid growth of cities has developed an increasing number of points of interest (POIs), e.g., restaurants, stores, hotels, providing us with more opportunities to experience life than ever before. Generally, people are willing to explore the city and neighborhood in their daily life, and decide "where to go" according to their personal interests w.r.t. the various choices of POIs. In the meantime, how to efficiently make a satisfying decision among the large number of POIs becomes a tough problem for a user, usually referred to as "choice paralysis" (Bensoussan and Fleisher 2012). POI recommendation is a task that aims to address such a problem by helping users filter out uninteresting POIs and reduce their decision making time (Ye et al. 2011).

With the development of location-based social networks (LBSNs), POI recommendation on LBSNs has recently attracted much attention as the latter provides an unprecedented opportunity to study human mobile behavior for POI

Table 1: Check-in Actions w.r.t. Content Information

Content Information	Facets of Check-in Actions
POI Properties	What is this POI about?
User Interests	Am I interested?
Sentiment Indications	How good is this POI?

recommendation in spatial, temporal, social, and content aspects (Gao and Liu 2014). Typical location-based social networking sites, e.g., Foursquare¹, Yelp², and Facebook Places³, allow users to "check in" at POIs with mobile devices like smartphones, and leave tips and share that experience with their online friends, resulting in a "W⁴" (i.e., **who**, **when**, **where**, and **what**) information layout, corresponding to four distinct information layers. While prior work of POI recommendation on LBSNs mostly focuses on investigating the spatial, temporal, and social layers in terms of spatial-temporal patterns (Gao et al. 2013; Cheng et al. 2013), social-spatial properties (Scellato et al. 2011; Gao, Tang, and Liu 2012), etc., the use of content information has not been systematically studied, thus missing the potential effect of content information in its recommendation efforts on LBSNs.

Content information on LBSNs could be related to a user's check-in action, providing a unique opportunity for POI recommendation. For example, by observing a POI's description as "vegetarian restaurant", we infer that the restaurant serves "vegetarian food" and users who check-in at this POI might be interested in the vegetarian diet. This is an example of *POI Properties*. By observing a user's comment on a Mexican restaurant discussing its spicy food, we observe the *User Interests* in spicy food. If the comment is actually a compliment, e.g., "Best spicy food ever!", we could infer both the user's *Sentiment Indications* and her interests. These three types of content information, i.e., *POI properties*, *User Interests*, and *Sentiment Indications*, are all related to a user's check-in actions and provide conceptual interpretations to three facets of her check-in actions, as listed in Table 1.

In recommender systems, user interests and target proper-

¹<http://foursquare.com>

²<http://www.yelp.com>

³<http://www.facebook.com/about/location>

ties are the two essential elements in capturing a user’s action (e.g., check-in) on a target (e.g., POI) for recommendation (Koren 2008), while user assessment has also been recognized as an important factor to gauge the check-in action for future recommendation (Sridharan 2001). Thus, content information on LBSNs provides a conceptual perspective to investigate users’ check-in behavior, which in turn constitutes the key factors of recommender systems, suggesting its potential for improving POI recommendation. In this work, we systematically study the three types of content information and propose a unified framework to model them for POI recommendation. Our contributions are summarized below.

- Study the relationship between users’ check-in behavior and content information on LBSNs in terms of POI properties, user interests, and sentiment indications.
- Incorporate the three types of content information into a unified framework for POI recommendation on LBSNs.
- Investigate the recommendation effort of each type of content information on a real-world LBSN dataset.

A Content-Aware POI Recommender System

In this section, we introduce a content-aware recommendation framework for POI recommendation on LBSNs.

Problem Statement

For ease of presentation, we use POI, venue, and location interchangeably in this work. Let $\mathbf{u} = \{u_1, u_2, \dots, u_M\}$ be the set of users and $\mathbf{v} = \{v_1, v_2, \dots, v_N\}$ be the set of venues, where M and N denote the numbers of users and POIs, respectively. $\mathbf{C} \in \mathbb{R}^{M \times N}$ is a check-in action matrix with each element C_{ij} representing the number of observed check-ins made by u_i at v_j . $\mathbf{S} \in \mathbb{R}^{M \times N}$ is a sentiment indication matrix extracted from users’ tips with $S_{ij} \in [-1, 1]$ representing the sentiment preference of u_i on v_j , where 1 and -1 represent the most positive and negative sentiment preference, respectively. Let $\mathbf{w} = \{w_1, w_2, \dots, w_T\}$ be the set of words, where T denotes the number of words. $\mathbf{A} \in \mathbb{R}^{M \times T}$ is a user-word matrix also obtained from users’ tips with A_{ij} as the frequency of w_j used in u_i ’s tips, representing user-interest content. $\mathbf{B} \in \mathbb{R}^{N \times T}$ is a POI-word matrix, with B_{ij} as the frequency of w_j appearing in v_i ’s description, representing POI-property content. The problem of content-aware POI recommendation can be formally described as,

Recommending POIs to a user based on observed check-in actions \mathbf{C} with the consideration of available content information, i.e., user sentiment indications \mathbf{S} , user-interest content \mathbf{A} , and POI-property content \mathbf{B} .

Exploring Content Information for POI Recommendation

Due to the user-driven check-in property (Gao and Liu 2014), the large-scale check-in data on LBSNs are usually very sparse. To solve large-scale recommendation problems with sparse data, matrix factorization is a state-of-the-art methodology that has been proven to be efficient and effective in movie recommendation (Koren 2008), item recommendation (Ma et al. 2008), and POI recommendation (Gao

et al. 2013). Thus, in this work, we leverage the content information on LBSNs with a low-rank matrix factorization method for POI recommendation. We first introduce a basic POI recommendation model based on matrix factorization without considering content information, then discuss how to incorporate the three types of content information into the basic model.

A Basic POI Recommendation Model Let $\mathbf{U} \in \mathbb{R}^{M \times K}$ be the users’ latent interests, $\mathbf{V} \in \mathbb{R}^{N \times K}$ be the POIs’ latent properties, and $\mathbf{H} \in \mathbb{R}^{K \times K}$ be the data-dependent dense matrix with $K \ll \min(M, N)$ being the number of latent factors. The basic POI recommendation model approximates u_i ’s latent interests in an unvisited v_j by solving the following optimization problem:

$$\min_{\mathbf{U}, \mathbf{H}, \mathbf{V}_{j \geq 0}} \frac{1}{2} \sum_i^m \sum_j^n \mathbf{W}_{ij} (\mathbf{C}_{ij} - \mathbf{U}_i \mathbf{H} \mathbf{V}_j^T)^2, \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{M \times N}$ is a check-in weighting matrix with $\mathbf{W}_{ij} = 1$ indicating that u_i has checked in at v_j , $\mathbf{W}_{ij} = 0$ otherwise.

The above recommendation model learns an optimal set of $\{\mathbf{U}, \mathbf{H}, \mathbf{V}\}$ whose product $\hat{\mathbf{C}} = \mathbf{U} \mathbf{H} \mathbf{V}^T$ is a non-sparse matrix which approximates the original \mathbf{C} . POI recommendation is then performed for each user based on the ranking among her unvisited POIs in $\hat{\mathbf{C}}$. Note that a non-negative constraint has been applied to \mathbf{U}_i , \mathbf{H} , and \mathbf{V}_j , respectively, as we consider that a user’s latent interests and a POI’s latent properties could have real-world explanations on LBSNs.

Modeling User Sentiment Indications As discussed above, \mathbf{W} is a weighting matrix applied to indicate the importance of check-ins, i.e., how likely the check-in action should be considered based on its observability, which is critical for improving recommendation performance (Pan and Scholz 2009). Previous work has discussed its potential effect when combined with information such as user reputation (Tang et al. 2013) or user activities (Li et al. 2010). This inspires us to investigate how to incorporate sentiment information for capturing check-in behavior.

Sentiment information is embedded in the tips on LBSNs that reflect users’ check-in experience. For example, if a user leaves a positive tip (or she likes it), the corresponding check-in is more important; otherwise it is less so. Thus, sentiment information can play a role as \mathbf{W} in Eq. (1) does in determining the importance of check-ins. To incorporate sentiment information, we propose a **sentiment-enhanced weighting scheme** as a function $\hat{\mathbf{W}} = f(\mathbf{W}, \mathbf{S})$ which assigns weights on check-ins based on the corresponding check-in observability and sentiment indications. We use $\hat{\mathbf{W}}$ to replace the original weighting matrix while the function $f(\cdot)$ should have the following properties:

- **Sentiment Consistency**
For an observed check-in action, a positive sentiment indication should increase its importance, while a negative sentiment indication should decrease its importance.
- **Sentiment Scaling**
To avoid over-weighting or under-weighting of sentiment information, the sentiment score in \mathbf{S} should be tuned to an appropriate scale before adopted for recommendation.

- **Non-Negativity**

The value in \mathbf{W} generated by $f(\cdot)$ should be non-negative according to the learning model in Eq. (1).

In this work, we empirically set $f(\cdot)$ as below, which works well in our recommendation model,

$$\hat{\mathbf{W}} = \mathbf{W} + \eta * \mathbf{S}, \quad \eta \in [0, 1], \quad (2)$$

where η is a scalar to control the weight from sentiment indications corresponding to the **Sentiment Scaling** property. Since observed check-ins have original weight $\mathbf{W}_{ij} = 1$, while the corresponding sentiment score $\mathbf{S}_{ij} \in [-1, 1]$, the hybrid score $\hat{\mathbf{W}}$ on an observed check-in is guaranteed to have the **Non-Negativity** property.

In Eq. (2), the importance of check-in actions is related to the corresponding sentiment score. A higher sentiment score \mathbf{S}_{ij} results in a greater value of $\hat{\mathbf{W}}_{ij}$, which forces $\mathbf{U}_i \mathbf{H} \mathbf{V}_j^T$ to tightly fit the check-in \mathbf{C}_{ij} while $\mathbf{U}_i \mathbf{H} \mathbf{V}_j^T$ will loosely approximate \mathbf{C}_{ij} when $\hat{\mathbf{W}}_{ij}$ is smaller (corresponding to a lower \mathbf{S}_{ij}). In an extreme case, when $\hat{\mathbf{W}}_{ij}$ is 0, the check-in action is not considered at all; thus, its likelihood of being recommended to other users is reduced. This is consistent to the user's sentiment indication, as $\hat{\mathbf{W}}_{ij} = 0$ happens when u_i presents the most negative score (-1) towards v_j (assuming $\eta=1$), which meets the **Sentiment Consistency** property.

Modeling User-Interest and POI-Property Content Besides user sentiment indications, user-interest content is also embedded in tips on LBSNs. Tips contain semantic words that reflect a user's interested topics regarding POIs, e.g., environment, taste, service, etc. It could help address the sparsity problem of check-in actions to a certain extent, as insufficient observation of check-in behavior can be compensated by the observed tipping behavior for inferring user interests. Thus, we propose the leveraging of information from tips to improve the learning of user latent interests, as shown below,

$$\min \frac{1}{2} \sum_i^M \sum_j^N (\mathbf{A}_{ij} - \mathbf{U}_i \mathbf{G}_j)^2, \quad (3)$$

where $\mathbf{G}_j \in \mathbb{R}^{K \times T}$ represents the word latent topics in user-interest content.

Similarly, information from POI-property content can also be leveraged to learn the POI latent properties, as shown below,

$$\min \frac{1}{2} \sum_i^M \sum_j^N (\mathbf{B}_{ij} - \mathbf{V}_i \hat{\mathbf{G}}_j)^2, \quad (4)$$

where $\hat{\mathbf{G}}_j \in \mathbb{R}^{K \times T}$ represents the word latent topics in POI-property content.

Both \mathbf{G}_j and $\hat{\mathbf{G}}_j$ represent word latent topics, where the former is in user context related to user-interest content, and the latter is in POI context related to POI-property content. Thus, we expect these two latent topics to be different but with certain overlaps, and propose a $\ell-1$ norm to capture such relationship,

$$\min \|\mathbf{G} - \hat{\mathbf{G}}\|_1 \quad (5)$$

where $\|\cdot\|_1$ is the $\ell-1$ regularization, with $\|X\|_1 = \sum_i \sum_j |X_{ij}|$.




Figure 1: Content-Aware POI Recommendation Framework

Content-Aware POI Recommendation Framework

Figure 1 illustrates our Content-Aware POI Recommendation Framework (CAPRF). Check-in action \mathbf{C} is directly related to sentiment indications \mathbf{S} , user interests \mathbf{U} , and POI properties \mathbf{V} , where the latter two are learned from the factorization of \mathbf{C} with the consideration of a data dependent matrix \mathbf{H} for model flexibility. User interests \mathbf{U} is also related to user-interest content \mathbf{A} , which represents tipping actions factorized to \mathbf{U} and word latent topics \mathbf{G} . POI properties \mathbf{V} is also related to POI-property content \mathbf{B} , which represents POI descriptions factorized to \mathbf{V} and word latent topics $\hat{\mathbf{G}}$. \mathbf{G} and $\hat{\mathbf{G}}$ are considered to have certain overlaps. The input of our framework is user check-in action \mathbf{C} , user sentiment indications \mathbf{S} , user-interest content \mathbf{A} , and POI-property content \mathbf{B} , and the output is \mathbf{U} , \mathbf{H} , and \mathbf{V} , whose product $\mathbf{U} \mathbf{H} \mathbf{V}^T$ is used for POI recommendation.

The proposed framework aims to solve the following optimization problem,

$$\begin{aligned} \min_{\mathbf{U}, \mathbf{H}, \mathbf{V} \geq 0} \mathcal{J} = & \frac{1}{2} \|\hat{\mathbf{W}} \odot (\mathbf{C} - \mathbf{U} \mathbf{H} \mathbf{V}^T)\|_F^2 + \frac{\lambda_1}{2} \|\mathbf{A} - \mathbf{U} \mathbf{G}\|_F^2 \\ & + \frac{\lambda_2}{2} \|\mathbf{B} - \mathbf{V}(\mathbf{G} - \mathbf{D})\|_F^2 + \delta \|\mathbf{D}\|_1 \\ & + \frac{\alpha}{2} (\|\mathbf{U}\|_F^2 + \|\mathbf{H}\|_F^2 + \|\mathbf{V}\|_F^2 + \|\mathbf{G}\|_F^2) \end{aligned} \quad (6)$$

where $\mathbf{D} = \mathbf{G} - \hat{\mathbf{G}}$ is proposed to facilitate the solving of $\ell-1$ regularization originally applied on both \mathbf{G} and $\hat{\mathbf{G}}$. λ_1 and λ_2 are introduced to control the weight of user-interest content and POI-property content. δ is a parameter to control the overlap between \mathbf{G}_j and $\hat{\mathbf{G}}$. The regularization terms $\|\mathbf{U}\|_F^2$, $\|\mathbf{H}\|_F^2$, $\|\mathbf{V}\|_F^2$, $\|\mathbf{G}\|_F^2$, and $\|\hat{\mathbf{G}}\|_F^2$ are used to avoid overfitting.

Parameter Estimation

Since there are multiple variables in the object function, we apply an alternative algorithm to find optimal solutions for the five variables \mathbf{U} , \mathbf{H} , \mathbf{V} , \mathbf{G} , and \mathbf{D} . The key idea is to minimize the objective function w.r.t. one variable while fixing the other variables, as similar to (Paatero and Tapper 1994). The algorithm will keep updating the variables until convergence or reaching the number of maximum iterations.

The objective function \mathcal{J} in Eq. (6) is differentiable at \mathbf{U} , \mathbf{H} , \mathbf{V} , and \mathbf{G} , the derivation of \mathcal{J} with respect to them are

$$\begin{aligned}\frac{\partial \mathcal{J}}{\partial \mathbf{U}} &= -(\hat{\mathbf{W}} \odot \hat{\mathbf{W}} \odot \mathbf{C})\mathbf{V}\mathbf{H}^\top + (\hat{\mathbf{W}} \odot \hat{\mathbf{W}} \odot (\mathbf{U}\mathbf{H}\mathbf{V}^\top))\mathbf{V}\mathbf{H}^\top \\ &\quad - \lambda_1 \mathbf{A}\mathbf{G}^\top + \lambda_1 \mathbf{U}\mathbf{G}\mathbf{G}^\top + \alpha \mathbf{U} \\ \frac{\partial \mathcal{J}}{\partial \mathbf{H}} &= -\mathbf{U}^\top (\hat{\mathbf{W}} \odot \hat{\mathbf{W}} \odot \mathbf{C})\mathbf{V} + \mathbf{U}^\top (\hat{\mathbf{W}} \odot \hat{\mathbf{W}} \odot (\mathbf{U}\mathbf{H}\mathbf{V}^\top))\mathbf{V} + \alpha \mathbf{H} \\ \frac{\partial \mathcal{J}}{\partial \mathbf{V}} &= -(\hat{\mathbf{W}}^\top \odot \hat{\mathbf{W}}^\top \odot \mathbf{C}^\top)\mathbf{U}\mathbf{H} + (\hat{\mathbf{W}}^\top \odot \hat{\mathbf{W}}^\top \odot (\mathbf{V}\mathbf{H}^\top \mathbf{U}^\top))\mathbf{U}\mathbf{H} \\ &\quad - \lambda_2 (\mathbf{B} - \mathbf{V}\mathbf{G} + \mathbf{V}\mathbf{D})(\mathbf{G}^\top - \mathbf{D}^\top) + \alpha \mathbf{V} \\ \frac{\partial \mathcal{J}}{\partial \mathbf{G}} &= -\lambda_1 \mathbf{U}^\top \mathbf{A} + \lambda_1 \mathbf{U}^\top \mathbf{U}\mathbf{G} - \lambda_2 \mathbf{V}^\top (\mathbf{B} - \mathbf{V}\mathbf{G} + \mathbf{V}\mathbf{D}) + \alpha \mathbf{G}\end{aligned}\quad (7)$$

The gradient descent optimization method is widely applied to update the above variables, and usually works well in recommender systems (Koren 2008). For the non-negative constraints on \mathbf{U} , \mathbf{H} , and \mathbf{V} , we applied projected strategy, which projects a negative parameter value to 0 in each iteration. The detailed updating rules are shown in Algorithm 1, where γ_u , γ_h , γ_v , and γ_g are learning steps that chosen to satisfy Goldstein Conditions (Jorge Nocedal 1999).

The optimization function regarding \mathbf{D} is non-smooth but convex, which can be formulated as a classical lasso problem. In this work, we use proximal gradient descent (Liu, Ji, and Ye 2009) to solve it.

Algorithm Analysis The detailed learning algorithm of our content-aware recommendation framework is shown in Algorithm 1. In lines 1-2, all the parameters are firstly initialized randomly. From lines 3 to 10, the algorithm iteratively updates \mathbf{U} , \mathbf{H} , \mathbf{V} , \mathbf{G} , and \mathbf{D} until convergence. The final output of this algorithm is $\hat{\mathbf{C}}$, which is the product of \mathbf{U} , \mathbf{H} , and \mathbf{V} . The overall time complexity of Algorithm 1 is (# of iterations)*($O(MKN) + O(NKT) + O(MKT)$).

Algorithm 1 Learning Algorithm of the Proposed Model

Input: user-POI check-in matrix \mathbf{C} , sentiment indication matrix \mathbf{S} , user-interest content \mathbf{A} , POI-property content \mathbf{B} , parameters $\{\eta, \lambda_1, \lambda_2, \delta, \alpha\}$

Output: approximated user-POI preference matrix $\hat{\mathbf{C}}$

- 1: Initialize \mathbf{U} , \mathbf{H} , \mathbf{V} , and \mathbf{G} randomly
 - 2: Set $\mathbf{W} = \text{sign}(\mathbf{C})$, $\hat{\mathbf{W}} = \mathbf{W} + \eta * \mathbf{S}$
 - 3: **while** Not Convergent **do**
 - 4: Calculate $\frac{\partial \mathcal{J}}{\partial \mathbf{U}}$, $\frac{\partial \mathcal{J}}{\partial \mathbf{H}}$, $\frac{\partial \mathcal{J}}{\partial \mathbf{V}}$, and $\frac{\partial \mathcal{J}}{\partial \mathbf{G}}$
 - 5: Update $\mathbf{U} \leftarrow \max(\mathbf{U} - \gamma_u \frac{\partial \mathcal{J}}{\partial \mathbf{U}}, 0)$
 - 6: Update $\mathbf{H} \leftarrow \max(\mathbf{H} - \gamma_h \frac{\partial \mathcal{J}}{\partial \mathbf{H}}, 0)$
 - 7: Update $\mathbf{V} \leftarrow \max(\mathbf{V} - \gamma_v \frac{\partial \mathcal{J}}{\partial \mathbf{V}}, 0)$
 - 8: Update $\mathbf{G} \leftarrow \mathbf{G} - \gamma_g \frac{\partial \mathcal{J}}{\partial \mathbf{G}}$
 - 9: Update \mathbf{D} with proximal gradient descent
 - 10: **end while**
 - 11: **return** $\hat{\mathbf{C}} = \mathbf{U}\mathbf{H}\mathbf{V}^\top$
-

Table 2: Statistical Information of the Dataset

No. of Users	4,287
No. of Check-ins	134,556
No. of POIs	5,878
No. of Tips	19,741
No. of Comments	56,718
Check-in Duration	May, 2008-Sep, 2013

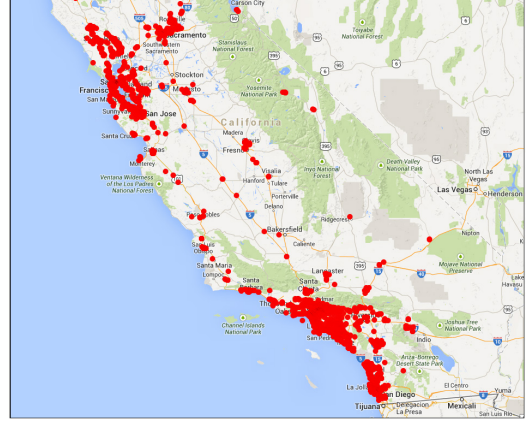


Figure 2: Check-in Distribution over California

Experiments

In this section, we evaluate the performance of our proposed framework **CAPRF** for POI recommendation. In particular, we evaluate the following: (1) how the proposed framework fares in comparison with state-of-the-art recommender systems; and (2) how different kinds of content information perform in the POI recommendation task. Before we delve into experiment details, we first discuss an LBSN dataset and evaluation metrics.

Foursquare Dataset

We choose Foursquare, one of the most popular location-based social networking sites, to study the content information on LBSNs. We collect users whose Foursquare profiles indicate their hometown as California state. We then obtain their corresponding check-in tweets with the same crawling strategy as proposed in (Scellato et al. 2011; Gao, Tang, and Liu 2012), and collect check-ins that happened in California state. Based on the venue id extracted from check-in tweets, we obtain the POI category through the ‘‘Venue API’’⁴ of Foursquare. We select check-ins happened at POIs of the ‘‘Food’’ category, which is the largest category among all the POIs in Foursquare. The POI-associated content (tags) is collected through the venue API as well, and user-generated content (tips) is collected through the ‘‘Tip API’’⁵. We consider users who have checked-in at least 2 distinct POIs. The statistics of the final dataset are listed in Table 2. Figure 2 presents the check-in distribution over California.

⁴<https://developer.foursquare.com/docs/venues/venues>

⁵<https://developer.foursquare.com/docs/users/tips>

Experimental Setup

The input of our framework is observed check-in action matrix \mathbf{C} and three types of content information \mathbf{S} , \mathbf{A} , and \mathbf{B} . We organize the check-in actions as a user-POI matrix \mathbf{C} . The check-in density of the matrix is 5.34×10^{-4} . A mapping function $\frac{1}{1+x^{-1}}$ is adopted to normalize \mathbf{C} , which has been proven to work well for POI recommendation in previous work (Gao et al. 2013).

We generate user sentiment indication matrix \mathbf{S} with an unsupervised sentiment classification method. For each tip, we remove stop words and employ a word-matching scheme to compute its sentiment score based on a sentiment lexicon. Sentiment polarity of a word is obtained from the pre-defined sentiment lexicon, i.e., -1 for negative and $+1$ for positive. The overall sentiment score of a tip is computed as the summation of sentiment scores of the words in the tip, and normalized to $[-1, 1]$ by taking the average on the tip length. We adopt the MPQA Subjectivity Lexicon⁶, which is a widely used manually labeled sentiment lexicon containing 2,718 positive and 4,902 negative words.

We select the common words of user-interest content and POI-property content, construct them as a user-word matrix \mathbf{A} and a POI-word matrix \mathbf{B} , with the matrix entry representing the frequency of a word used by corresponding user/POI. The total number of common words is 1,810.

For each individual user in the check-in matrix, we randomly mark off 20% of all POIs that he has checked-in for testing. The rest of the observed user-POI pairs are used as training data for POI recommendation. The random selection is conducted 5 times individually, and we report the average results. Since only the observed check-in actions (corresponding to $\mathbf{W}_{ij} = 1$) are considered in Eq. (1), following the standard strategy of solving one-class CF problems (Pan and Scholz 2009; Pan et al. 2008), we sample 10% of unobserved check-ins from the training matrix, deem them as the check-in frequency of 0 and set their corresponding \mathbf{W}_{ij} to 1. The same strategy is also performed on baseline methods.

We use *precision@N* and *recall@N* as our evaluation metrics. In our experiment, N is set to be 5 and 10. All the parameters in this paper are set through cross-validation. For the proposed method, the experimental results use $d=20$ dimensions to represent the latent features, the parameters $\{\eta, \lambda_1, \lambda_2, \delta, \alpha\}$ are set to $\{0.3, 0.1, 0.1, 0.8, 0.1\}$.

Performance Evaluation

In this section, we compare our POI recommendation framework with five existing state-of-the-art methods: (1) User-Based Collaborative Filtering (**UCF**) (Zhou et al. 2012); (2) Probabilistic Matrix Factorization (**PMF**) (Salakhutdinov and Mnih 2007); (3) Non-negative Matrix Factorization (**NMF**) (Lee, Seung, and others 1999); (4) Spatial Topic Location Recommender (**STLR**) (Hu and Ester 2013); and (5) Sentiment-Enhanced Location Recommender (**SELR**) (Dingqi Yang and Wang 2013). The first three are classical collaborative filtering approaches for general recommendation problems without considering content effects, while the last two consider user-interest content only and

Table 3: Performance Comparison

Methods	Precision		Recall	
	P@5	P@10	R@5	R@10
UCF	0.0083	0.0077	0.0117	0.0216
PMF	0.0114	0.0104	0.0160	0.0292
NMF	0.0126	0.0111	0.0177	0.0310
STLR	0.0173	0.0150	0.0243	0.0422
SELR	0.0134	0.0121	0.0188	0.0340
CAPRF	0.0186	0.0169	0.0261	0.0474

sentiment indications only, respectively. In addition, we further investigate the recommendation efforts of each type of content information with our proposed model, which gives the interpretation on the effect of POI-property content.

Table 3 reports the comparison results of **CAPRF** with the proposed baseline methods. The results precipitate several observations, which we summarize below:

- UCF performs the worst among all the approaches. Data sparseness is a possible reason for this performance. Due to the low density of the check-in matrix, the user-based collaborative filtering approach fails to accurately recommend POIs and performs worse than matrix factorization approaches, i.e., PMF and NMF, which leverage the low-rank approximation of user check-in preferences.
- SELR and STLR perform better than UCF, PMF, and NMF, suggesting the importance of sentiment information and user-interest content. Furthermore, the better performance of STLR over SELR indicates that user-interest content seems to be more effective than sentiment information for POI recommendation. We will further discuss this in the next section.

Among all the approaches, CAPRF performs the best, suggesting the importance of content information on LBSNs for POI recommendation. It is worth noting that precision and recall in our experiments are not high. As suggested in (Ye, Liu, and Lee 2012), the effectiveness of recommender systems with sparse datasets is usually very low. For example, the reported top 5 precision is 5% over a dataset with 8.02×10^{-3} density (Ye, Liu, and Lee 2012). Similar performance can also be observed in other POI recommendation work (Liu et al. 2013; Yin et al. 2013; Hu and Ester 2013). Therefore, the low precision obtained in our experiment is reasonable. In this paper, we focus on **comparing algorithms' relative performance instead of their absolute performance**.

Evaluation of Three Types of Content Information

In this section, we discuss the recommendation efforts of different types of content information on LBSNs, i.e., Sentiment Indications (**SI**), User-Interest Content (**UIC**), and POI-Property Content (**PPC**). We propose to consider their different combinations by setting the corresponding parameter η (for sentiment indications), λ_1 (for user-interest content), and λ_2 (for POI-property content). For each parameter, if the corresponding type of content is considered, we set it to the optimal value; otherwise, 0. Since δ relates to

⁶http://mpqa.cs.pitt.edu/lexicons/subj_lexicon

Table 4: Recommendation Effect of Different Types of Content Information

Information	SI	UIC	PPC	P@5	R@5	P@10	R@10
NONE	×	×	×	0.0126	0.0177	0.0111	0.0310
SI	√	×	×	0.0130	0.0182	0.0115	0.0323
UIC	×	√	×	0.0164	0.0230	0.0148	0.0416
PPC	×	×	√	0.0154	0.0217	0.0138	0.0387
UIC+SI	√	√	×	0.0172	0.0242	0.0154	0.0433
PPC+SI	√	×	√	0.0157	0.0221	0.0141	0.0396
UIC+PPC	×	√	√	0.0178	0.0250	0.0162	0.0456
SI+UIC+PPC	√	√	√	0.0186	0.0261	0.0169	0.0474

both **UIC** and **PPC**, we set it to the optimal value when both types of content information have been considered, and 0 if only one or none of them has been considered. Table 4 lists the comparison results. We use \sqrt to indicate that the corresponding content information is used, and \times otherwise.

Sentiment information is helpful in improving the POI recommendation performance. It consistently improves the performance based on existing content information. However, its recommendation effect is not as great as user-interest content and POI-property content. One possible reason could be that sentiment information is quite noisy in user-generated content, while the non-perfect lexicon-based sentiment classification approach exacerbates the capturing of essential user attitude content.

User-interest content presents more recommendation effect than the other two types of content information. One possible reason could be that user-interest content contains more information than POI-property content regarding a user’s interests, which is more helpful to capture user preferences. The combination of all three types of content information, i.e., CAPRF, has the best performance among all the other methods. This indicates the complementary effect among the three types of content information. According to Table 1, these information constitutes the key factors of POI recommender systems.

Related Work

POI recommendation, also referred to as location recommendation, has been recognized as an essential task on recommender systems. It was firstly studied on GPS trajectory data (Zheng et al. 2009; Zheng and Xie 2011). With the development of LBSNs, users are able to check-in at real-world POIs which generates abundant spatial, temporal, social, and content information. Due to the strong correlations between geographical distance and social connections discovered in previous work (Cheng et al. 2011; Cho, Myers, and Leskovec 2011; Scellato et al. 2011; Noulas et al. 2011; Zhang et al. 2012; Gao, Tang, and Liu 2012; Long and Joshi 2013; Chang and Sun 2011), current work on POI recommendation on LBSNs mainly focuses on leveraging the geographical and social properties to improve recommendation effectiveness. Ye et al. (Ye, Yin, and Lee 2010) introduced POI recommendation into LBSNs and investigated the geographical influence (Ye et al. 2011) and social influence (Ye, Liu, and Lee 2012) for POI recommendation. Cheng (Cheng

et al. 2012) investigated the geographical and social influence through a multi-center Gaussian model.

Temporal information has also attracted much attention from researchers. Gao et al. (Gao et al. 2013) investigated the temporal cyclic patterns of check-ins in terms of temporal non-uniformness and temporal consecutiveness. Yuan et al. (Yuan et al. 2013) incorporated both temporal cyclic information and geographical information for time-aware POI recommendation. Cheng et al. (Cheng et al. 2013) introduced the task of successive personalized POI recommendation in LBSNs with a matrix factorization method.

Most recently, researchers started to explore the content information on LBSNs for POI recommendation. Yang et al. (Dingqi Yang and Wang 2013) introduced sentiment information and reported its better performance over state-of-the-art approaches. Hu et al. (Hu and Ester 2013) investigated the user-interest content from Twitter and Yelp. Liu et al. (Liu et al. 2013; Liu and Xiong 2013) studied the effect of POI-associated tags with an aggregated LDA model. Yin et al. (Yin et al. 2013) investigated both personal interest and local preference on LBSNs and EBSNs. All these work focuses on one type of content information without considering the other two and their correlations.

Conclusion and Future Work

In this work, we systematically study the content information on LBSNs for POI recommendation. We investigate various types of content information on LBSNs in terms of sentiment indications, user interests, and POI properties. We model them under a unified POI recommendation framework. Our experimental results demonstrate the significance of content information in explaining user behavior and improve POI recommendation performance on LBSNs. In the future, we will continue studying the content information, especially the effect of sentiment indications, such as user opinions, on different facets of a POI. Furthermore, it would be interesting to investigate the recommendation effect of content information compared to other information, such as spatial, temporal, or social information.

Acknowledgments

This material is based upon work supported by, or in part by, the U.S. Army Research Office (ARO) under contract/grant number 025071, and the Office of Naval Research (ONR) under grant number N000141010091.

References

- Bensoussan, B. E., and Fleisher, C. S. 2012. *Analysis Without Paralysis: 12 Tools to Make Better Strategic Decisions*. FT Press.
- Chang, J., and Sun, E. 2011. Location 3: How users share and respond to location-based data on social networking sites. *Proceedings of the ICWSM*.
- Cheng, Z.; Caverlee, J.; Lee, K.; and Sui, D. 2011. Exploring millions of footprints in location sharing services. In *Proceedings of the ICWSM*.
- Cheng, C.; Yang, H.; King, I.; and Lyu, M. 2012. Fused matrix factorization with geographical and social influence in location-based social networks. *AAAI, Canada*.
- Cheng, C.; Yang, H.; Lyu, M. R.; and King, I. 2013. Where you like to go next: Successive point-of-interest recommendation. In *Proceedings of the IJCAI*, 2605–2611.
- Cho, E.; Myers, S.; and Leskovec, J. 2011. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the SIGKDD*, 1082–1090.
- Dingqi Yang, Daqing Zhang, Z. Y., and Wang, Z. 2013. A sentiment-enhanced personalized location recommendation system. *ACM Hypertext*.
- Gao, H., and Liu, H. 2014. Data analysis on location-based social networks. In *Mobile Social Networking*. Springer. 165–194.
- Gao, H.; Tang, J.; Hu, X.; and Liu, H. 2013. Exploring temporal effects for location recommendation on location-based social networks. In *Proceedings of the RecSys*, 93–100.
- Gao, H.; Tang, J.; and Liu, H. 2012. Exploring social-historical ties on location-based social networks. In *Proceedings of the ICWSM*.
- Hu, B., and Ester, M. 2013. Spatial topic modeling in online social media for location recommendation. In *Proceedings of the RecSys*, 25–32.
- Jorge Nocedal, S. W. 1999. *Numerical Optimization*. Springer.
- Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceeding of the SIGKDD*, 426–434.
- Lee, D.; Seung, H.; et al. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401(6755):788–791.
- Li, Y.; Hu, J.; Zhai, C.; and Chen, Y. 2010. Improving one-class collaborative filtering by incorporating rich user information. In *Proceedings of the CIKM*, 959–968. ACM.
- Liu, B., and Xiong, H. 2013. Point-of-interest recommendation in location based social networks with topic and location awareness. *Proc. of SDM* 396–404.
- Liu, B.; Fu, Y.; Yao, Z.; and Xiong, H. 2013. Learning geographical preferences for point-of-interest recommendation. In *Proceedings of the SIGKDD*, 1043–1051. ACM.
- Liu, J.; Ji, S.; and Ye, J. 2009. Multi-task feature learning via efficient $l_2, 1$ -norm minimization. In *Proceedings of the UAI*, 339–348. AUAI Press.
- Long, X., and Joshi, J. 2013. A hits-based poi recommendation algorithm for location-based social networks. In *Proceedings of the ASONAM*, 642–647. ACM.
- Ma, H.; Yang, H.; Lyu, M.; and King, I. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the CIKM*, 931–940. ACM.
- Noulas, A.; Scellato, S.; Mascolo, C.; and Pontil, M. 2011. An empirical study of geographic user activity patterns in foursquare. *Proceeding of the ICWSM*.
- Paatero, P., and Tapper, U. 1994. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* (2).
- Pan, R., and Scholz, M. 2009. Mind the gaps: weighting the unknown in large-scale one-class collaborative filtering. In *Proceedings of the SIGKDD*, 667–676. ACM.
- Pan, R.; Zhou, Y.; Cao, B.; Liu, N. N.; Lukose, R.; Scholz, M.; and Yang, Q. 2008. One-class collaborative filtering. In *Proceeding of the ICDM*, 502–511. IEEE.
- Salakhutdinov, R., and Mnih, A. 2007. Probabilistic matrix factorization. In *NIPS*, volume 1, 2–1.
- Scellato, S.; Noulas, A.; Lambiotte, R.; and Mascolo, C. 2011. Socio-spatial properties of online location-based social networks. *Proceeding of the ICWSM* 11.
- Sridharan, S. 2001. *Usability and Reliability of the User Action Framework: A Theoretical Foundation for Usability Engineering Activities*. Ph.D. Dissertation, Citeseer.
- Tang, J.; Hu, X.; Gao, H.; and Liu, H. 2013. Exploiting local and global social context for recommendation. In *Proceedings of the IJCAI*, 2712–2718. AAAI Press.
- Ye, M.; Yin, P.; Lee, W.; and Lee, D. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *Proceeding of the SIGIR*, 325–334.
- Ye, M.; Liu, X.; and Lee, W. 2012. Exploring social influence for recommendation - a probabilistic generative approach. In *Proceeding of the SIGIR*, 325–334.
- Ye, M.; Yin, P.; and Lee, W. 2010. Location recommendation for location-based social networks. In *Proceedings of the SIGSPATIAL*, 458–461.
- Yin, H.; Sun, Y.; Cui, B.; Hu, Z.; and Chen, L. 2013. Lcars: a location-content-aware recommender system. In *Proceedings of the SIGKDD*, 221–229. ACM.
- Yuan, Q.; Cong, G.; Ma, Z.; Sun, A.; and Thalmann, N. M. 2013. Time-aware point-of-interest recommendation. In *Proceedings of the SIGIR*, 363–372. ACM.
- Zhang, C.; Shou, L.; Chen, K.; Chen, G.; and Bei, Y. 2012. Evaluating geo-social influence in location-based social networks. In *Proceedings of the CIKM*, 1442–1451. ACM.
- Zheng, Y., and Xie, X. 2011. Learning travel recommendations from user-generated gps traces. *TIST* 2(1):2.
- Zheng, Y.; Zhang, L.; Xie, X.; and Ma, W. 2009. Mining interesting locations and travel sequences from gps trajectories. In *WWW*, 791–800. ACM.
- Zhou, D.; Wang, B.; Rahimi, S.; and Wang, X. 2012. A study of recommending locations on location-based social network by collaborative filtering. *Proceeding of the AAI*.