Exploiting Geo-Spatial Preference for Personalized Expert Recommendation

Haokai Lu
Texas A&M University
hlu@cse.tamu.edu

James Caverlee
Texas A&M University
caverlee@cse.tamu.edu

ABSTRACT

Experts are important for providing reliable and authoritative information and opinion, as well as for improving online reviews and services. While considerable previous research has focused on finding topical experts with broad appeal – e.g., top Java developers, best lawyers in Texas – we tackle the problem of *personalized expert recommendation*, to identify experts who have special personal appeal and importance to users. One of the key insights motivating our approach is to leverage the geo-spatial preferences of users and the variation of these preferences across different regions, topics, and social communities. Through a fine-grained GPS-tagged social media trace, we characterize these geo-spatial preferences for personalized experts, and integrate these preferences into a matrix factorization-based personalized expert recommender. Through extensive experiments, we find that the proposed approach can improve the quality of recommendation by 24% in precision compared to several baselines. We also find that users' geo-spatial preference of expertise and their underlying social communities can ameliorate the cold start problem by more than 20% in precision and recall.

Categories and Subject Descriptors

H.2.8 [**Database Applications**]: Data Mining

Keywords

expert recommendation; geospatial preference; GPS-tagged social media

1. INTRODUCTION

Finding and recommending *experts* is a critical component for many important tasks. For example, the quality of movie recommenders can be improved by biasing the underlying models toward the opinions of experts [1]. Making sense of mobile and social information streams such as the Facebook newsfeed and the Twitter stream can be improved by focusing on content contributed by experts. Along these lines, companies like Google and Yelp are actively soliciting *expert reviewers* to improve the coverage and reliability of their services [8]. More generally and in contrast to search engines and question-answer systems, experts can provide ongoing help for evolving and ill-specified needs, as well as personalized access to knowledge and experience that only experts possess.

Indeed, there has been considerable effort toward expert finding and recommendation, e.g., [2, 3, 6, 7, 14, 17, 20, 22]. These efforts have typically sought to identify topical experts with broad appeal, e.g., the top Java developer in an enterprise, the best lawyer in Texas. However, there is a research gap in our understanding of both (i) identifying *personal experts*, that is experts who are of significance and importance to me, but perhaps not viewed so more broadly. For example, I may be interested in the expert opinions of nearby local foodies, but less interested in the opinions of globally popular celebrity chefs; and (ii) how spatial preference for personally-valuable expertise varies across topics, across regions, and based on different underlying social communities. For example, technologists in Houston, TX may be more interested in the opinions of experts in nearby Austin and in more distant Silicon Valley, but less so in the opinions of experts from New York. Similarly, the reach of experts may vary by location, so that tech experts from Silicon Valley have a larger footprint than do experts from other regions.

Hence, in this paper, we are interested to study the problem of *personalized expert recommendation* by integrating the geo-spatial preferences of users and the variation of these preferences across different regions, topics, and social communities. These geo-spatial preferences are increasingly being revealed through the fine-grained geo-spatial footprints of Instagram, Foursquare, and Twitter, among other mobile location sharing platforms. Concretely, we opportunistically leverage a collection of GPS-tagged Twitter users and their relationships in Twitter lists, a form of crowd-sourced knowledge whereby user A may label user B with a descriptor (like "technology"). In isolation these lists allow a user to organize a personal Twitter stream; in aggregate, the many labels applied to a target user in many lists can provide a crowdsourced expertise profile of the target user. Specifically, we propose and evaluate a matrix factorization-based personalized expert recommender that leverages three key factors:

- *Region-based locality*, reflecting the variation in spatial preference from region to region. For example, Figure 1b and Figure 1c shows that the preference of users for food experts varies greatly based on the location of the user (in essence, local users prefer local foodies). How can these regional differences be captured and incorporated into a personalized expert recommender?
- *Topic-based locality*, reflecting the variation in spatial preference across different topics. For example, Figure 1b and Figure 1e demonstrate that spatial preference is much less local for the topic technology than for food. How can this topical variation be integrated into a personalized expert recommender?

RecSys'15, September 16–20, 2015, Vienna, Austria.

DOI: http://dx.doi.org/10.1145/2792838.2800189.

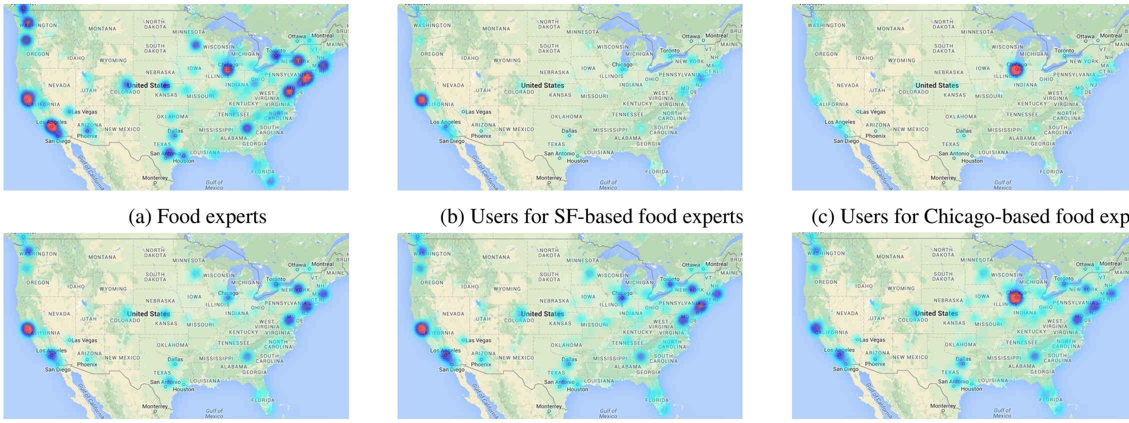

(a) Food experts (b) Users for SF-based food experts (c) Users for Chicago-based food experts

(d) Tech experts (e) Users for SF-based tech experts (f) Users for Chicago-based tech experts

Figure 1: Spatial distribution of experts (a,d) and for the users who have listed experts (b,c,e,f) based on geo-tagged Twitter lists.

- *Social-based locality*, reflecting the social connections between users and experts. For example, are users who are connected in an underlying social network more "similar" in their preferences for experts? Are experts who are more tightly coupled in the underlying social network preferred by the same set of users?

Through extensive experimental validation, we find that each of these factors – region, topic, and social-based locality – improves the quality of personalized expert recommendation. And together, the proposed model achieves around 24% improvement in precision and 21% improvement in recall versus both a collaborative filtering and a baseline matrix factorization based recommender. Furthermore, we also find that the proposed approach can ameliorate the cold start problem when users have few experts on their lists, leading to more than 20% improvement over the baseline in precision and recall.

2. RELATED WORK

Many previous works [2, 3, 14, 22] have focused on finding general topic experts in many domains (e.g., enterprise corporate, email networks), with a recent emphasis on social media and microblogging sites [7, 17, 20]. Weng et al. [20] proposed a PageRank-based approach to find topic experts by taking advantage of both topical similarity between users and social link structure. Pal and Counts [17] introduced a probabilistic clustering followed by a within-cluster Gaussian ranking framework to find topic authorities using nodal and topical features on Twitter. Ghosh et al. [7] proposed and built a system called Cognos to find topic experts by relying on Twitter Lists (though not with any geo-spatial information, as in this work). Recently, Cheng et al. [6] addressed the problem of identifying local experts on Twitter. Our work extends on these prior efforts by focusing on *personalized* experts and in investigating the impact of region-based, topic-based, and social community-based locality on personalized expert recommendation.

With the rapid growth of location-based social networks in recent years, many applications [5, 9, 12, 13, 15, 21] have started to take advantage of geo footprints such as point-of-interest (POI) recommendation on social networks. Ye et al. [21] explored the spatial clustering phenomenon and proposed a unified POI recommendation framework combining user preference, geographical influence, and social influence. Cheng et al. [5] proposed a multi-center Gaussian model to model user's check-in behavior, which is used as input for a generalized matrix factorization framework. Liu et al. [13] proposed a geographical probabilistic factor analysis framework, which jointly models the effect of geographical distance, user preference, POI popularity and user mobility. Another different application that utilizes geographical footprints is the rating prediction problem in Yelp [9], where Hu et al. observed weak positive correlation between a business's ratings and its neighbor's ratings, and used this observation to improve rating predictions. In contrast, we are focused on preferences for experts, rather than on particular POIs or venues.

3. PERSONALIZED EXPERT RECOMMENDATION: OVERVIEW

In this section, we introduce the problem of personalized expert recommendation and outline our core approach.

3.1 Problem Statement

We assume there exists a set of users $U = \{u_1, u_2, ..., u_N\}$, where N is the total number of users. From this set U, there are a number of recognized experts denoted as $E = \{e_1, e_2..., e_M\}$, where M is the total number of experts. Each user has a preference over some of these experts, expressed as a *personalized expertise list*. For example, Alice may prefer Beth to Candace in the topic of "Java programming", but have no opinion on Doug. We then define the problem of *personalized expert recommendation* as: Given a user u_i, identify the top-n personally relevant experts to u_i. That is, can we further identify experts Eva and Frank that are of personal interest to Alice?

3.2 Recommendation by Matrix Factorization

We tackle *personalized expert recommendation* using latent factor matrix factorization [10]. We assume there is a factor p_i associated with each user u_i and a factor q_j associated with each expert e_j. The model defines a rating score between the list and the expert, denoted as y_{ij}, and factors the score into a latent space through p_i and q_j as follows:

$$y_{ij} = p_i^T q_j + b_j \tag{1}$$

Through this factorization, we can think of q_j as the latent properties for expert e_j, p_i as the latent preference of user u_i and b_j as the popularity bias for e_j. However, unlike the standard recommendation task, we do not have a rating score for each expert on the lists. Instead, we only have the implicit feedback for a list, which assumes a user prefers an expert who is already on the list to an expert who is not. Accordingly, the learning objective should be based on the pair-wise ranking between experts. In recommendations when only implicit feedback is available, the one-class collaborative filtering approach [4, 18, 23] can be used for learning a rank order among items. Similar efforts have been targeted at tag recommendation [19], tweet recommendation [4], and event-based groups [11, 23]. Here, we adapt the Bayesian Personalized Ranking (BPR) criterion proposed by Rendle et al. in [18] to our problem.

Formally, for a user u_i, an expert e_k and an expert e_h, suppose u_i puts e_k on the list while not e_h, we denote this pair as $e_k^{u_i} \succeq e_h^{u_i}$, and the likelihood for this preference under BPR can be written as:

$$p(e_k^{u_i} \succeq e_h^{u_i}) = \sigma(y_{ik} - y_{ih}) \text{ where } \sigma(x) = \frac{1}{1+e^{-x}}$$

Therefore, the likelihood for all users could be written as:

$$p(R|\Theta) = \prod_{e_k^{u_i} \in \mathcal{P}^{u_i}, e_h^{u_i} \in \mathcal{N}^{u_i}, u_i \in U} p(e_k^{u_i} \succeq e_h^{u_i})$$

where R is the set of all preference pairs, Θ is the set of all parameters, $\mathcal{P}^{u_i}$ is the set of experts included on u_i's list and $\mathcal{N}^{u_i}$ is the set of absent experts for u_i. If Θ has a prior density $p(\Theta)$, we can derive a bayesian version of the likelihood, where the prior is used to prevent the overfitting of the parameters as a form of regularization. Thus, the posterior log-likelihood to maximize is

$$p(R|\Theta) = \sum_{e_k^{u_i} \in \mathcal{P}^{u_i}, e_h^{u_i} \in \mathcal{N}^{u_i}, u_i \in U} ln(\sigma(y_{ik} - y_{ih})) - \text{regularization}$$

which can be learned through stochastic gradient descent (SGD) by iterating each of the preference pairs and updating the corresponding parameters.

4. REGION, TOPIC, AND SOCIAL-BASED LOCALITY

While promising, the baseline matrix factorization approach ignores the geo-spatial preferences of users and the variation of these preferences across different topics, regions, and social communities (as suggested by Figure 1's intuitive support for these notions). Hence, we turn in this section to demonstrating how these factors manifest in real-world Twitter-based data and how each of these factors can be incorporated into a new personalized expert recommendation matrix factorization framework.

4.1 Data and Metrics

We begin by highlighting the data used here and two statistical measures – expert entropy and expert spread – to characterize region and topic-based locality. We then turn to the social properties of the dataset to demonstrate social-based locality.

Data. We use the geo-tagged Twitter lists collected in [6]. In total, there are about 12 million crowd-generated lists and 14 million geo-tagged listings, where a geo-tagged listing indicates a direct link from a list creator to an expert where both of their geo-locations are known. That is, each user $u_i \in U$ is associated with geographical coordinates $coord_{u_i}$. Furthermore, for each list, there exist associated labels that list creators use to indicate the topic of that list. In the following analysis, we selected lists which include the most frequent unigram labels indicating typical topics as follows: news, music, tech, sports, celebs, and food. Additionally, we randomly sampled lists which include any unigram occurring more than 200 times in list labels. We denote this randomly sampled list data as "general". Furthermore, we excluded experts who have only occurred in one list and also excluded lists which includes only one expert. After filtering, we have the geo-tagged Twitter list data statistics shown in Table 1. In the following sections, we refer to list creators as *users* and list members as *experts*.

Table 1: Geo-tagged Twitter list data.

topic	# of lists	# of experts	# of listings	sparsity(%)
news	35,539	20,295	287,321	0.04
music	17,945	7,896	160,286	0.11
sports	16,018	5,395	139,838	0.16
food	10,476	5,485	96,661	0.17
celebs	9,783	4,090	104,004	0.26
tech	13,046	10,760	125,178	0.26
general	30,000	36,217	289,528	0.03

Metrics. We discretize the continental US surface with a 1° by 1° geodesic grid to map the coordinates to discrete regions.[1] Formally, we have a total number of K grids, which we call regions. We denote K regions as $R = \{r_i | i = 1, 2, ..., K\}$, to which each coordinate inside the US can be mapped. Furthermore, we assume for an expert e, there are totally n_e users who put e on their lists. Among them, we let U^e be the set of users for expert e, and $U^e_{r_i}$ be the set of users from the region r_i. Thus, the probability of expert e's user from the region r_i can be defined as $p^e_{r_i} = \frac{|U^e_{r_i}|}{\sum_{r_i \in R} |U^e_{r_i}|}$. With these preliminaries, we quantity the geographical characteristics of expertise with:

Expert entropy. The *expert entropy* is defined as

$$H(e) = -\sum_{r_i \in R} p^e_{r_i} log(p^e_{r_i})$$

This measure indicates the degree of randomness in spatial distribution of the users for an expert. It ranges from 0 when all users for the expert are only from one region, to $logK$ when user's distribution is uniform across all regions. Thus, it implicitly reflects the level of an expert's recognizability across the entire country.

Expert spread. While entropy provides insights into the spatial distribution of users, it lacks explicit consideration for the distance between a user and an expert. Hence, we define another measure called *expert spread* as follows:

$$S(e) = Median_{u_i \in U^e}(d(coord_e, coord_{u_i}))$$

where d is the distance between two locations, computed with Haversine function to account for the shape of the earth. The expert spread indicates how far a typical user is from an expert, thus can be considered as the localness of an expert.

4.2 Region-Based Locality

In Figure 1b and 1c, we observed that food experts from San Francisco and Chicago are preferred by users nearby. How does this observation manifest according to our statistical measures? To that end, we select experts from the following cities: San Francisco (SF), New York (NY), Chicago, Houston, Denver and Seattle. We first show the average expert entropy for these cities with respect to different topics in Table 2. As can be observed from the table: (i) Experts from different geo-locations have different levels of recognizability across the country; and (ii) Generally, experts from SF and NY are popular in more regions than those from other geo-locations, indicating that SF and NY have a greater impact on expertise curation for users on Twitter.

In Figure 2a, we examine expert spread for these cities. We can see that generally, experts from different geo-locations have different levels of locality, with experts from Chicago and SF having the

[1] 1° by 1° is approximately 70 miles by 50 miles at latitude 40°. We also tested a finer mesh of 0.1° by 0.1°, which gave quantitatively similar results.

Table 2: Average expert entropy for different cities.

topic	SF	NY	Houston	Chicago	Seattle	Denver
news	**2.461**	2.342	2.021	1.950	1.836	1.884
music	**2.514**	2.386	1.946	1.996	2.105	2.162
sports	2.518	**2.703**	1.956	2.281	2.217	2.060
food	1.689	**2.105**	1.315	1.172	1.439	1.327
celebs	**3.274**	2.777	3.013	2.781	2.950	2.842
tech	2.323	**2.400**	2.262	2.249	2.098	1.917
general	**1.954**	1.932	1.645	1.610	1.606	1.585

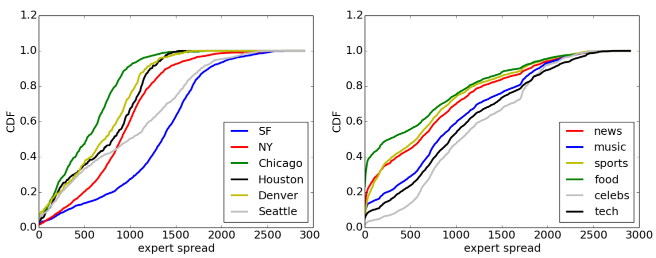

(a) (b)

Figure 2: CDF of expert spread (a) for different cities and (b) for different topics.

smallest and largest expert spread. This indicates that compared to other cities, Chicago has the most local influence on expertise curation while SF reaches the farthest. Combined with the observations from Table 2, we conclude that *experts from different regions may have different levels of locality*, i.e., some may reach a wider geographical scope but others may be only locally popular.

Integrating region-based locality. Since the observed region-based locality reflects collective opinion, how can we integrate it into personalized expert recommendation? That is, if we know the geo-location of a user, can we recommend experts who are popular around the user's geo-location? As can be observed in Figure 1, an expert's popularity is not necessarily linear in the distance between user and expert; rather, it is often in the form of "clusters", i.e., experts may be popular in one region but not in other regions. Thus, we introduce the concept of "regional popularity", where we parameterize the popularity of each expert by regions, with the regional popularity to be learned from training.

Concretely, we assume the geographical space is partitioned into K regions. For an expert e_j and a region r_i, we assume there is a popularity parameter s_{ij} associated with r_i and e_j. This parameter is used to capture the degree of popularity that expert e_j receives in region r_i. Thus, if we have a total of M experts, the popularity parameters constitute the matrix S of dimension K by M, which represents regional popularity for all experts. Each column $s_{\cdot j}$ represents the popularity e_j receives in all regions. Then, given a user u_i, the popularity e_j receives at the region where u_i is from is denoted as $s_{c(u_i)j}$, where $c(x)$ is a function mapping a user to its region. We use $s_{c_i j}$ instead of $s_{c(u_i)j}$ for convenience. By integrating the matrix S to the original matrix factorization, we have:

$$y_{ij} = p_i^T q_j + s_{c_i j} \tag{2}$$

We denote Equation 2 as the Geo-Enhanced factorization (GEF). Note that GEF is reduced to Equation 1 when $K = 1$. The GEF approach has the advantage over the baseline matrix factorization of explicitly capturing and learning expert regional popularity.

Table 3: Average expert entropy and expert spread (miles) when $CDF = 0.5$ for different topics.

topic	food	news	tech	sports	music	celebs
entropy	1.661	2.048	2.235	2.247	2.267	**2.868**
spread	290	630	950	580	830	**1060**

4.3 Topic-Based Locality

In the previous analysis of expertise, we observed that expert entropy can be impacted by geo-locations (see Table 2). Additionally, this table also implies that expert entropy can be impacted by the choice of topic. To further observe the geo-spatial distribution of expertise for different topics, we list the average expert entropy for the six sample topics in Table 3. As we can see, celebs has the largest entropy, which indicates that users interested in celebrity are most widely spread across the country; while food has the smallest entropy, indicating that users interested in food experts are most concentrated in certain regions. This is intuitively reasonable since a celebrity is very likely to have a better chance of being known in the whole country than a food expert from a certain location.

In Figure 2b, we show the cumulative density function against expert spread for different topics. We can see that, for a fixed spread value, the topic food gives the largest cumulative probability, indicating that users interested in food are closest to the experts; while users interested in celebrity are farthest. We also show the expert spread when the CDF is 0.5 in Table 3. We can see that the topics with increasing expert spread are ordered as: food < sports < news < music < tech < celebs, with food having the smallest expert spread of 290 miles, and celebs having the largest expert spread of 1060 miles, which is almost half the distance from the west coast to east coast of the US. Combined with the previous observation on expert entropy, we can conclude that the topic food is the most local among all, with users mostly concentrated in local regions of experts, while the topic celebs is the least local, with users scattered across the country. In another word, users interested in food tend to select food expert nearby, while users interested in celebrity do not have such geographical constraints, and users interested in other topics fall in between.

Overall, we can conclude that *expert's regional popularity can vary by topic*; in other words, users may have different regional preference for experts because of their topic interests.

Integrating topic-based locality. Now that we have observed that topic locality can influence user's preference for experts, it is important that user's interests should be aligned with the interests of the experts to be recommended. Since each Twitter list is labeled with certain keywords, we can aggregate all of the labels for an expert in all lists he appears. As a result, an expert e_j has a description d_{e_j} consisting of the aggregated labels. We then introduce a user latent topic factor t_{u_i}, representing u_i's topical preference, and expert topic factor $\bar{t}_{e_j}$, representing the topical property of e_j. Thus, the inner product $t_{u_i}^T \bar{t}_{e_j}$ indicates an affinity score of user u_i and expert e_j with respect to topic. Here, $\bar{t}_{e_j}$ is treated as known through d_{e_j}, and t_{u_i} is treated as unknown to be learned. The reason to model in this way is that labels for lists often have only one term, e.g., lists with one term label "food" occupy about 60% percent of total lists with any "food" in its labels. But often, a list is very focused on finer aspects of a topic. For example, a list labeled with "food" may include many "wine" experts, implying that we should also consider expert candidates labeled with "wine". By making t_{u_i} unknown, we are forcing the model to learn topic aspects of a user from those of experts she selected. For convenience, we use t_i instead of t_{u_i} and $\bar{t}_j$ instead of $\bar{t}_{e_j}$ afterwards. Thus, our

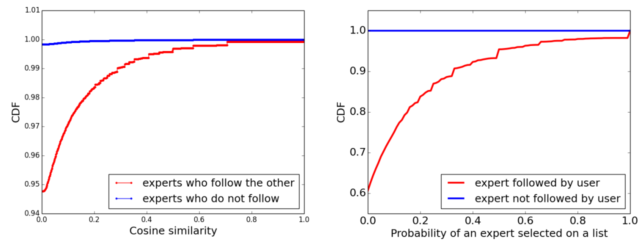

(a) (b)

Figure 3: (a): CDF for similarity between experts; (b) CDF for the probability of an expert on a list.

Topic-Enhanced factorization (TEF) can be written as:

$$y_{ij} = p_i^T q_j + t_i^T \bar{t}_j \tag{3}$$

Here, we treat each label as a dimension of $\bar{t}_j$. Through the explicit handling of each user's topic aspects, it is expected that user's interests are aligned with the interests of the experts to be recommended.

Fusion of region and topic-based locality. Naturally, we can integrate both region and topic-based locality into the model. We adopt a linear model for the integration of Equation 2 and Equation 3, resulting in our Geo-Topic Enhanced factorization (GTEF):

$$y_{ij} = p_i^T q_j + s_{c_i j} + t_i^T \bar{t}_j \tag{4}$$

The intuition is when we know the region of a user and her topic aspects (by looking at the labels of her selected experts), we can recommend an expert both topically and geographically relevant.

4.4 Social-Based Locality

In addition to the modeling of region and topic based locality, we are also interested to explore if social connections among users and experts can improve expert recommendation. Our intuition is that (i) people who are connected by social ties have a higher probability to have similar interests; (ii) people who are socially related may have a higher probability to select who they follow as experts.

As evidence of social-based locality, in Figure 3a, we compare the similarity of experts for two cases: (i) when one expert follows the other; (ii) when no tie exists between two experts. Here, similarity of experts is defined as the cosine similarity computed by viewing each expert as a vector of all users, with each element being a value indicating whether the expert is listed by the user or not. Thus, a large similarity of two experts indicates that they often occur on the same list. We can observe that experts who follow the other generally have a larger similarity. We also compare the probability of an expert selected on a list in Figure 3b for two cases: (i) when experts are followed by the user; and (ii) when experts are not followed by the user. We can see that the chance for an expert to be listed by a user is boosted significantly when that expert is already followed by the user. Based on these observations, we individually model three kinds of social relationships:

User-user relationship. In this case, the following relationship is from a list creator to another list creator. When one user follows another, we assume that their preference is more similar to each other than those who do not. In terms of modeling, we adopt the approach of regulating their latent factors as in [16]. Formally, suppose there are user u_i and u_j, assume $\mathcal{F}_i^u$ is the set of users u_i follows, the social regularization incurred by user user relationship can be written as:

$$\sum_{i=1}^{N} \sum_{f \in \mathcal{F}_i^u} w(u_i, u_f) ||p_i - p_f||^2$$

where $w(u_i, u_f)$ represents the similarity between u_i and u_f. Thus, if u_i and u_f is more similar, the latent preference factor p_i and p_f is also closer. Here, we use cosine similarity of users as the weighting scheme. The cosine similarity of users is computed by viewing each user as a vector of all experts, with each element taking a value — 1 if the expert is on the list, 0 if not.

Expert-expert relationship. In this case, the following relationship is from an expert to another expert. Using a similar approach as in the previous case, we regulate their latent factors so that experts following the other have similar latent factors. Formally, assume $\mathcal{F}_i^e$ is the set of experts e_i follows, the social regularization incurred by expert expert relationship can be written as:

$$\sum_{j=1}^{M} \sum_{f \in \mathcal{F}_j^e} w(e_j, e_f) ||q_j - q_f||^2$$

where $w(e_j, e_f)$ represents the similarity between e_j and e_f. Here, we also use cosine similarity of experts as the weighting scheme.

User-expert relationship. In this case, the following relationship is from a user to an expert. Unlike the previous two kinds of social ties, this relationship links two different entities, and so the regularization approach is ill-suited here. Instead, we explicitly model these relationships with a bias term added to Equation 4 as follows:

$$y_{ij} = p_i^T q_j + s_{c_i j} + t_i^T \bar{t}_j + \theta_i b_{ij} \tag{5}$$

where b_{ij} takes a boolean value 1 if u_i follows e_j, and 0 if not. θ_i is a weighting parameter to be learned. Thus, by adding a personalized bias term for each user, the model can take advantage of the following ties between user and expert.

4.5 Model Training

Combining the social regularization and Equation 5, the final objective function to maximize can be written as:

$$\sum_{e_k^{u_i} \in \mathcal{P}^{u_i}, e_h^{u_i} \in \mathcal{N}^{u_i}, u_i \in U} ln(\frac{1}{1 + e^{-(y_{ik} - y_{ih})}})$$
$$- \frac{\beta_1}{2} \sum_{i=1}^{N} \sum_{f \in \mathcal{F}_i^u} w(u_i, u_f) ||p_i - p_f||^2$$
$$- \frac{\beta_2}{2} \sum_{j=1}^{M} \sum_{f \in \mathcal{F}_j^e} w(e_j, e_f) ||q_j - q_f||^2 - regularization$$

where L_2-norm regularization is adopted, with β_1 and β_2 as the corresponding regularization parameters. In summary, the parameter set Θ to be learned through SGD is $\{p_i, q_j, t_i, \theta_i, s_{.j}\}$. For each iteration of SGD, we need to sample a user u_i, and from u_i's list, an expert e_k. Due to the large size of absent experts for each user, we also need to sample the set $\mathcal{N}^{u_i}$. Here, we adopt the strategy of random sampling. Then, for each sampled triplet $< l_i, e_k, e_h >$, we update each parameter value by taking a step along its gradient ascending:

$$\Theta^{t+1} = \Theta^t + \epsilon \frac{\partial L_{ikh}}{\partial \Theta}$$

where L_{ikh} is the posterior log-likelihood for the triplet $< l_i, e_k, e_h >$, and ϵ is the step size.

Update of S. For expert e_k and e_h, the corresponding parameters to update are $s_{.k}$ and $s_{.h}$. If the region of u_i is c_i, then the parameter $s_{.k}$ and $s_{.h}$ can be updated with

$$\frac{\partial L_{ikh}}{\partial s_{jk}} = -I(j = c_i)\hat{e} + \beta s_{c_i k}, \frac{\partial L_{ikh}}{\partial s_{jh}} = I(j = c_i)\hat{e} + \beta s_{c_i h}$$

where $\hat{e} = \frac{e^{-(y_{ik}-y_{ih})}}{1+e^{-(y_{ik}-y_{ih})}}$, $I(j = c_i)$ is a Kronecker delta function that gives value 1 if and only if $j = c_i$, for $j = 1, ..., K$, and β is a regularization parameter.

Update of t_i and θ_i. Similarly, we have the gradient for t_i and θ_i as follows:

$$\frac{\partial L_{ikh}}{\partial t_i} = -\hat{e}(\bar{t}_k - \bar{t}_h) + \beta t_i, \frac{\partial L_{ikh}}{\partial \theta_i} = -\hat{e}(b_{ik} - b_{ih}) + \beta \theta_i$$

Update of p_i, q_k and q_h. Since p and q are socially regularized, we have the following socially regularized gradients:

$$\frac{\partial L_{ikh}}{\partial p_i} = -\hat{e}(p_k - p_h) + \beta p_i + \beta_1 \sum_{f \in \mathcal{F}_i^u} w(u_i, u_f)(p_i - p_f)$$

$$\frac{\partial L_{ikh}}{\partial q_k} = -\hat{e} p_i + \beta q_k + \beta_2 \sum_{f \in \mathcal{F}_i^e} w(e_k, e_f)(q_k - q_f)$$

The gradient for q_h can be obtained similarly as q_k.

5. EXPERIMENTAL EVALUATION

In this section, we report on experiments to evaluate the proposed Geo-Topic Enhanced Factorization with Social ties (GTEF-S) for personalized expert recommendation. Specifically, we seek answers to the following questions:

- How well does the proposed method perform compared to alternative baselines? Does region, topic and social-based locality give improvement individually, and if they do, do they complement each other?
- How well does it perform in cold-start situation, i.e., for users who have very few experts on their lists?
- Does the number of regions affect performance? If so, how?

5.1 Data Preparation and Experimental Setup

For evaluation, we randomly partition experts for a user into 50% for training and 50% for testing. To determine the number of negative experts for each user, we experimented with $\{50, 100, 150, 200, 250\}$ and selected 150 for a tradeoff between accuracy and computational efficiency. For latent factor dimension, we empirically select 20 for all methods. For regularization parameters β, β_1 and β_2, we use cross-validation for tuning and select 0.02, 0.01 and 0.015, respectively. For gradient step, we initialize it with the step size 0.025, and decrease it to its 98% after each pass throughout all triples. This strategy is shown to be effective in reducing the number of iterations for the method to converge [9].

In the modeling of region locality, it is assumed that the continental US has been partitioned into K regions. Instead of using a gridding approach, we resort to k-means clustering to obtain the partitions by clustering the geo-locations of the entire set of users U. We choose a clustering approach based on Euclidean distance because the geo-spatial distribution of users exhibits a clustering effect, as shown in Figure 1, and can be satisfactorily captured by k-means clustering. In section 5.5, we evaluate the effect of the number of regions K. For other experiments, we select K to be 80.

For evaluation metrics, we adopt Precision@k (Prec@k) and Recall@k (Rec@k). Prec@k represents the percentage of correctly recommended experts out of the top k recommendations, while Rec@k represents what percentage of experts can emerge in the top k recommendations. Formally, if we define $Test(u)$ as the set of experts selected by user u and $Reco(u)$ as the set of top k recommended experts, we have

$$Prec@k = \frac{1}{N} \sum_{i=1}^{N} \frac{|Test(u_i) \cap Reco(u_i)|}{k}$$

$$Rec@k = \frac{1}{N} \sum_{i=1}^{N} \frac{|Test(u_i) \cap Reco(u_i)|}{|Test(u_i)|}$$

In our experiments, we evaluate k at 5, 10 and 15.

5.2 Baselines

We consider the following baselines:

- Expert Popularity (EP). In this baseline, we recommend experts for each user by ranking experts according to the number times each expert is listed by users.
- User-based Collaborative Filtering (UCF). Collaborative filtering method can be used to discover user's implicit preference by aggregating similar users. Formally, let a_{ij} take a boolean value, where $a_{ij} = 1$ represents expert e_j is selected by u_i, while $a_{ij} = 0$ means the opposite. Thus, according to UCF, the prediction score $\bar{c}_{ij}$ of u_i selecting e_j can be obtained by $\bar{c}_{ij} = \frac{\sum_k w_{ik} \cdot c_{kj}}{\sum_k w_{ik}}$, where w_{ik} is computed with cosine similarity. We then rank the candidate experts according to $\bar{c}_{ij}$ and select the top k experts for recommendation. We select the number of neighbors for each user to be 100.
- MF. This is the basic pair-wise latent factor model shown in Equation 1 trained by BPR.
- GEF. This model only considers region-based locality manifested through users' geographical footprints, shown in Equation 2.
- TEF. This model only considers topic-based locality manifested through experts' labels, shown in Equation 3.
- GTEF. This model is the fusion of GEF and TEF, considering both region and topic-based locality as shown in Equation 4.
- Social MF. This model considers three different kinds of social ties. If the model only considers user user relationship, it is denoted as MF-S1; if the model only considers expert expert relationship, it is denoted as MF-S2; and if the model only considers user expert relationship, it is denoted as MF-S3. We denote the model as MF-S if it considers all three kinds of social ties.

5.3 Comparison with Baselines

How well does the proposed method compare to alternative approaches? To answer this question, we first show the performance comparison in Figure 4, where we report Prec@k and Rec@k for all topics. As we can see, overall, the proposed GTEF-S generally gives the best performance for different k. Specifically, it gives an average improvement of 24.6% over the best of EP, UCF and MF for precision, and 21.3% for recall. GTEF-S generally performs better than either GTEF or MF-S, indicating the superiority of enhanced pair-wise matrix factorization by considering region, topic and social-based locality, and that these three factors are able to complement each other.

Comparison for MF, GEF, TEF and GTEF. By comparing these methods, we can examine if the explicit modeling of region and topic-based locality can provide any improvement. In Figure 4, we can see that GEF, TEF and GTEF perform consistently better than MF. Specifically, GEF gives an average improvement of 3.73% for precision and 3.43% for recall over MF for all datasets. This indicates that the introduction of the regional popularity matrix S for modeling expert's regional popularity can help distinguish regionally popular experts if we know the geo-location of the user.

Furthermore, we can see that TEF also performs consistently better than MF, specifically, giving an average improvement of 3.91% for precision and 3.16% for recall. This indicates that modeling user topic factor through the labels of experts can help find experts with similar topic aspect.

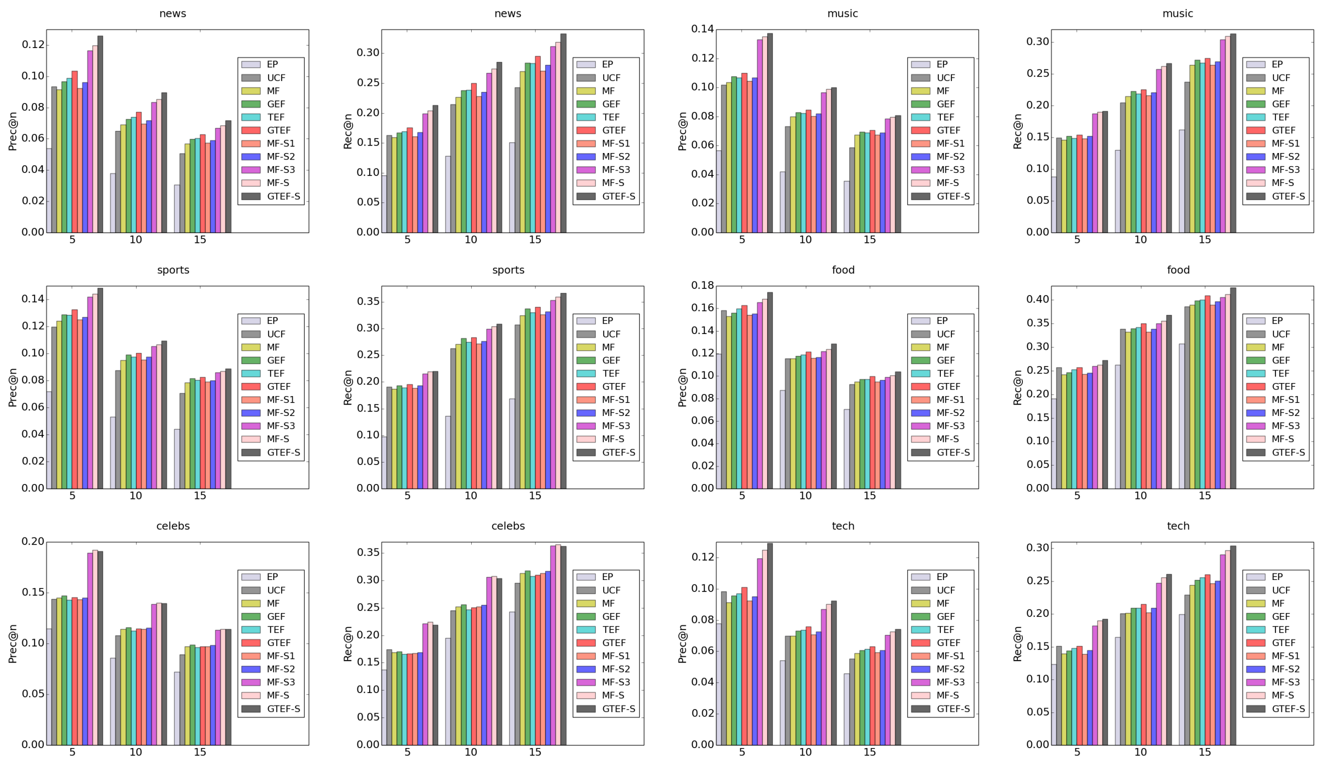

Figure 4: Evaluating personalized expert recommendation: Precision and Recall at 5, 10, and 15 for six different topics across 11 approaches.

Finally, we can see that GTEF gives the best performance among all (an average improvement of 7.35% for precision and 6.28% for recall over MF). These improvements are very close to the additive improvements of both GEF and TEF, thus indicating geographical influence is complementary to topic influence in modeling expert's regional popularity, and that they should be considered together for recommendation. Note that for topic celebs, the improvement of GEF (1.44% for precision) is not as good as those for other topics. This is probably because celebrities are heavily concentrated in the region of Los Angeles, and since its expert entropy is very high, it would not be very useful to model expert regional popularity. Also, TEF performs slightly worse than MF. Upon further examining the labels for experts and user's topic preference factors, we found that the labeling information for experts is scarce and most experts are only labeled with "celebs", without finer topic aspects.

Comparison for MF-S1, MF-S2, MF-S3 and MF-S. By comparing these methods, we can examine if the modeling of social-based locality can help recommend experts. Specifically, we have explored three kinds of relations: user following user (S1), expert following expert (S2) and user following expert (S3). From Figure 4, we can see that MF-S1 gives only slightly better performance than MF (0.43% for precision and 0.35% for recall), indicating that social ties between users barely provide additional information for recommending experts. For MF-S2, we see that it provides descent improvement over MF (5.02% for precision and 4.92% for recall), confirming that if an expert is similar to another expert, i.e., if they often co-occur on other lists, it is likely that the other expert can be recommended to this user. For MF-S3, it is shown to be rather effective. On average, it gives an improvement of 16.6% for precision and 14.4% for recall over MF, which confirms that if a user is already following this expert, it is very likely that this user will include this expert on the list. MF-S, modeling the previous three social relationships together, gives the best improvement of all (21.4% for precision and 18.8% for recall), indicating three kinds of social ties complement each other.

5.4 Recommendation for Cold-Start Lists

Previously, we found that the introduction of geographical, topical and social influence in pair-wise MF can improve expert recommendation. In this section, we examine how the proposed methods perform in the cold-start situation. When there is only limited number of experts on a list, there is little positive feedback for training, making it hard to obtain accurate latent factors of users' preferences. In consideration of this, we perform experiments to investigate the recommendation performance of the proposed methods for lists with few experts. Specifically, we select only lists which have fewer than 3 experts on the list to examine the performance. In Figure 5, we report the prec@5 and rec@5 for the method MF, GTEF, MF-S and GTEF-S. As we can see, GTEF, MF-S and GTEF-S consistently give better performance than MF for all topics, with GTEF-S showing the best improvement on average (23.7% for precision and 22.3% for recall). This indicates that the knowledge of user's region, topic preference and social relations can help relieve the cold-start problem. Also, MF-S gives better performance than GTEF, indicating that social relation is a stronger signal than user's geo-topic preference. Additionally, note that GTEF-S brings the best improvement over MF-S for news and food. This implies that modeling region and topic-based locality works best when users demonstrate strong regional preference (see Table 3).

5.5 Effect of Number of Regions

In this section, we study the effect of the number of regions K chosen to cluster the geographical coordinates of users. To that end, we select the number of regions from the set $\{10, 20, 40, 60, 80, 100\}$,

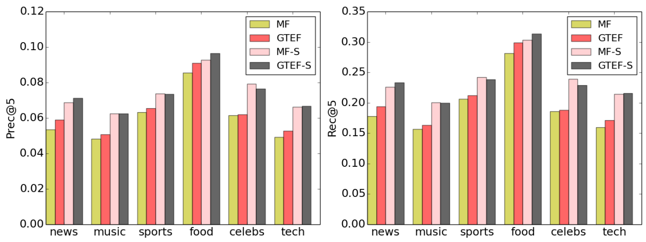

Figure 5: Comparing recommenders for cold start lists.

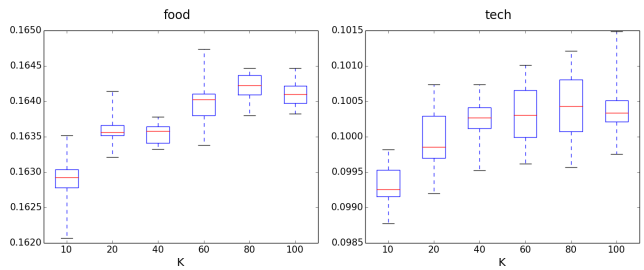

Figure 6: Effect of number of regions.

and run GTEF with each number for ten random initializations for topic food and tech (we ignore plots for other topics since they show similar trends). In Figure 6, we show how prec@5 changes with K. We can see that as the number of regions increases, the precision also generally increases, although the value of K varies for two topics when the performance reaches saturation. Specifically, for tech, the precision reaches almost the best at a smaller K, while for food, the precision gradually increases and reaches the best at a larger K. This can be explained by the observation from the previous analysis about topic locality. Specifically, topics such as food and news are relatively more local, i.e., experts of these topics are listed by local people more often. Also, experts of these topics are usually concentrated in many regions across the country, as shown in Figure 1a. As a result, a finer clustering of regions would separate two close regions. For example, it would separate the region of NY and Washington D.C. in Figure 1a, so that a user interested in food in NY can be recommended with popular food experts from NY instead of popular food experts from Washington D.C. On the other hand, topics such as tech and celebs are less local, and considering most of these experts are concentrated in fewer regions, it is not necessary to use a finer clustering.

6. CONCLUSION

In this paper, we tackled the problem of personalized expert recommendation in GPS-enabled social media. Specifically, we investigated the geo-spatial preferences of users and the variation of these preferences across different regions, topics and social communities. We proposed a matrix factorization-based personalized expert recommender that leverages region, topic and social-based locality. Through experimental evaluation over a Twitter list dataset, we found that the proposed approach achieves more than 20% in precision and recall and can ameliorate the cold start problem compared to several baselines. This confirmed users' geo-spatial preference of expertise and their underlying social communities have great potential for personalized expert recommendation.

7. ACKNOWLEDGMENTS

This work was supported in part by NSF grant IIS-1149383 and Google Research Award. Any opinions, findings and conclusions or recommendations expressed in this material are the author(s) and do not necessarily reflect those of the sponsors.

8. REFERENCES

[1] X. Amatriain, N. Lathia, J. M. Pujol, H. Kwak, and N. Oliver. The wisdom of the few: A collaborative filtering approach based on expert opinions from the web. In *SIGIR*, 2009.

[2] K. Balog, L. Azzopardi, and M. De Rijke. Formal models for expert finding in enterprise corpora. In *SIGIR*, 2006.

[3] C. Campbell, P. Maglio, A. Cozzi, and B. Dom. Expertise identification using email communications. In *CIKM*, 2003.

[4] K. Chen, T. Chen, G. Zheng, Q. Jin, E. Yao, and Y. Yu. Collaborative personalized tweet recommendation. In *SIGIR*, 2012.

[5] C. Cheng, H. Yang, I. King, and M. R. Lyu. Fused matrix factorization with geographical and social influence in location-based social networks. In *AAAI*, 2012.

[6] Z. Cheng, J. Caverlee, H. Barthwal, and V. Bachani. Who is the barbecue king of texas?: A geo-spatial approach to finding local experts on twitter. In *SIGIR*, 2014.

[7] S. Ghosh, N. Sharma, F. Benevenuto, N. Ganguly, and K. Gummadi. Cognos: Crowdsourcing search for topic experts in microblogs. In *SIGIR*, 2012.

[8] Google. Overview of local guides, April 2015.

[9] L. Hu, A. Sun, and Y. Liu. Your neighbors affect your ratings: On geographical neighborhood influence to rating prediction. In *SIGIR*, 2014.

[10] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 2009.

[11] A. Krohn-Grimberghe, L. Drumond, C. Freudenthaler, and L. Schmidt-Thieme. Multi-relational matrix factorization using bayesian personalized ranking for social network data. In *WSDM*, 2012.

[12] D. Lian, C. Zhao, X. Xie, G. Sun, E. Chen, and Y. Rui. Geomf: joint geographical modeling and matrix factorization for point-of-interest recommendation. In *SIGKDD*, 2014.

[13] B. Liu, Y. Fu, Z. Yao, and H. Xiong. Learning geographical preferences for point-of-interest recommendation. In *SIGKDD*, 2013.

[14] X. Liu, B. Croft, and M. Koll. Finding experts in community based question-answering services. In *CIKM*, 2005.

[15] Y. Liu, W. Wei, A. Sun, and C. Miao. Exploiting geographical neighborhood characteristics for location recommendation. In *CIKM*, 2014.

[16] H. Ma, D. Zhou, C. Liu, M. Lyu, and I. King. Recommender systems with social regularization. In *WSDM*, 2011.

[17] A. Pal and S. Counts. Identifying topical authorities in microblogs. In *WSDM*, 2011.

[18] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, 2009.

[19] S. Rendle and L. Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *WSDM*, 2010.

[20] J. Weng, E.P. Lim, J. Jiang, and Q. He. Twitterrank: Finding topic-sensitive influential twitterers. In *WSDM*, 2010.

[21] M. Ye, P. Yin, W.C. Lee, and D.L. Lee. Exploiting geographical influence for collaborative point-of-interest recommendation. In *SIGIR*, 2011.

[22] J. Zhang, M. S. Ackerman, and L. Adamic. Expertise networks in online communities: structure and algorithms. In *WWW*, 2007.

[23] W. Zhang, J. Wang, and W. Feng. Combining latent factor model with location features for event-based group recommendation. In *SIGKDD*, 2013.