Dynamic Item Recommendation by Topic Modeling for Social Networks

Sang Su Lee
Computer Science Department
University of Southern California
Los Angeles, CA 90089
Email: sangsl@usc.edu

Tagyoung Chung
Computer Science Department
University of Rochester
Rochester, NY 14627
Email: chung@cs.rochester.edu

Dennis McLeod
Computer Science Department
University of Southern California
Los Angeles, CA 90089
Email: mcleod@usc.edu

***Abstract*—The need to identify an approach that recommends items that match users' preferences within social networks has grown in tandem with the increasing number of items appearing within these networks. This research presents a novel technique for item recommendation within social networks that matches user and group interests over time. Users often tag items in social networks with words and phrases that reflect their preferred "vocabulary." As such, these tags provide succinct descriptions of the resource; implicitly reveal user preferences, and, as the tag vocabulary of users tends to change over time, reflect the dynamics of user preferences. Based on evaluation of user and group interests over time, we present a recommendation system employing a modified latent Dirichlet allocation (LDA) model in which users and tags associated with an item are represented and clustered by topics, and the topic-based representation is combined with the item's timestamp to show time-based topic distribution. By representing users via topics, the model can cluster users to reveal the group interests. Based on this model, we developed a recommendation system that reflects user as well as group interests in a dynamic manner that accounts for time, allowing it to perform in a manner superior to that of static recommendation systems in terms of precision rate.**

***Index Terms*—Web mining, Tagging, Recommender systems, Information analysis, Social network services.**

I. Introduction

Within social networks, the practice of *tagging*, the process of creating and using a *tag*, a relevant term that is associated with a unit of information, has become common. Users create and share their own content, such as blog posts, photographs, and videos, and then may create and use several tags to describe their content. A recommendation system can be considered a specific type of information filtering (IF) technique [1] that attempts to present information resources (e.g., images, videos, music, and URLs) that are likely to be of interest to users. Typically, a recommendation system compares a user's characteristics against some reference and attempts to predict how the user would rate an item that he or she has not yet rated. The characteristics of users may include those concerning an information item (the content-based approach) or the user's social environment (the collaborative-filtering [2] approach).

The need for recommendation systems has increased in tandem with the great increase in the number of information resources and the consequent difficulty of identifying relevant resources within social networks. Reacting to this need, the research into tag recommendation systems based on users' previous tag usages has recently been extended to research into item recommendation systems using tags. As described above, a tag describes the characteristics of an item, and may represent the characteristics of users who are associated with the item. For this reason, tags are useful indicators of characteristics in recommendation systems, particularly in cases in which it is difficult to quantitatively retrieve the characteristics of items (e.g., video, audio, or images) from the contents of the items.

Although we hypothesize that tag distribution changes over time, current research into the use of tagging in item recommendation has not considered changes in interests over time. For example, some groups of users who had been interested in football last fall may be interested in basketball this spring. In this case, the users' tag vocabulary contained more tags related to football in fall but more tags related to basketball in spring. Thus, at a specific point of time (e.g., March), it would be more desirable for a recommendation system to recommend items related to basketball to users.

To address this consideration, we developed a system that uses the process of dynamic adjustment and includes tags with similar concepts and interests to recommend items with greater precision. We modeled our system after the *latent Dirichlet allocation* (LDA) model [3], a generative model that approximates the generation of items in terms of latent topics. The LDA model considers an item a mixture of various latent topics, and chooses tags in the item according to the topics. We extended the LDA model in order to model users and tags over time by representing users and tags as mixtures of topics and by reflecting the time-based distribution of topics. To reflect changes in interests over time, we introduced the concept of a time-based similarity weight. Generally, a recommendation system suggests a new item to a user after determining which of the user's item is most similar to the new item. If the similarity between the new item and the user's item reaches or exceeds a threshold value, the system recommends the new item to the user. If the tag distribution differs over time, the similarity metric that determines the distribution change can better determine the similarities between items, thus allowing the system to make better recommendations.

Our work makes the following meaningful contributions. First, our recommendation system considers solely tags and not the items themselves to make recommendations. By doing so, the system remains simple in terms of types of attributes while, as we demonstrate, performing well, suggesting that tags provide accurate summaries of contents. Second, our system reflects changes in interests over time and makes use of change as a factor in recommendations. The remainder of the paper is organized as follows. Section 2 discusses related work, Section 3 describes our approach, Section 4 presents and evaluates our results, and Section 5 concludes our work and presents extensions for further research.

II. Related Work

This section examines three different areas of research that are related to our work. First, it examines work that focuses on recommendation systems which use tags as auxiliary information for traditional content-based or rating-based recommendation system. Next, Tag-based recommendation systems that only use tag as primary information for recommendation are explained.

Zhen et al. [4] proposed a novel framework to integrate tagging information into the CF procedure. In a similar manner, Sen et al. [5] extended the capability of the current movie recommendation systems by using tags as an indicator of user preferences. To recommend movies, the authors inferred user preferences for tags from user movie interaction such as movie ratings and clicks or users' tagging behavior, then used the inferred tag preferences to make movie recommendations. De Gemmis et al. [6] investigated whether folksonomies might be valuable sources of information regarding user interests, and might contribute to a strategy that enables a content-based recommender to infer user interests by applying machine-learning techniques on both the "official" item descriptions provided by a publisher and on the tags that users adopt to freely annotate relevant items.

In social networks, rating information is rare. Therefore, other research has focused on proposing recommendation systems which only use tags and do not use rating or other information for the generalized use of the recommendation system. Guan et al. [7] proposed a graph-based representation learning algorithm for this purpose according to which the users, tags, and documents are represented in the same semantic space in which two related objects are close to each other. Siersdorfer et al. [8] proposed a formal model to characterize users, items, and annotations within social networks to fulfill their goal of constructing a social recommendation system that predicts the utility of items, users, or groups based on the tagging vocabulary of a given user by implementing a LDA model. Guo et al. [9] proposed a recommendation system based on a probabilistic generative model for tagging. They introduced a modified LDA model, which is used to cluster the tags and users, to generate user as well as group interest information from the LDA model, and employed that information to recommend items to users.

A considerable number of research regarding tag-based recommendation system is done. The current research, however, does not focus on temporal aspect for recommendation. In our approach, we use temporal aspect for the better recommendation result by using tags and not using rating for the generalized approach.

III. Approach

Our approach for dynamic item recommendation has four parts: preprocessing a dataset from social networks; topic modeling from the preprocessed dataset; time-based similarity weight calculation; and recommendation. A dataset in our approach consists of three entities: an *item*, which is a resource preferred by a user; a *user*, who is an individual who prefers an item; and a *tag*, which is an entity annotated to the item to describe the item. The relationship between items and tags is *many-to-many*, as is the relationship between users and tags. To design our experiments, we extracted each user's items and associated tags. From the extracted data, we created a dataset in which each item's features are its tags. We then passed our dataset through the steps of pruning irrelevant tags, grouping users and documents by topic through a variant of the LDA model, and learning similarity weights over time. Then, we designed the system to recommend new testing items to each user based on the similarity between weights over time and between new testing items and his or her own items.

A. Tag Pruning

In our experiment, we collect the items and associated tags for each user, and represent each item as a vector of tags with which the item is annotated. Whereas some tags provide useful information for creating recommendations, other tags are too general or too specific to be useful. As it would be computationally difficult to use all the tags in the following step, we reduced the number of tags by determining their *term frequency-inverse document frequency* (tf-idf) weight [10], a statistical measure used to determine the importance of a term to a document in a collection or corpus.

In our dataset, we consider each tag a term and each user a document. After calculating the tf-idf weights of favorite tags for each user, we eliminate those tags whose weights fell below a certain threshold, leaving us with the tags that best describe a user's preferences.

B. LDA-based topic modeling over time

Using tf-idf weights for pruning, we effectively reduce the dimension of the problem by decreasing the number of tags. Nevertheless, the feature space remained too large. Fortunately, we could further reduce the dimension of the problem by the use of dimensionality reduction techniques, specifically the use of the LDA model [3] for dimensionality reduction. In the LDA model, which is widely used for topic discovery based on the occurrences of words (tags) in documents (items), an item is considered a bag of particular tags and represented by mixtures over latent topics, with each latent topic being characterized by a fixed conditional distribution over tags.

Using the LDA model is thus regarded as a form of topic modeling because of the topic representation for which it provides. Specifically, the LDA model assumes that all tags of all items are generated by randomly chosen latent topics. The following section introduces the basic concepts of the LDA model.

1) Extension of the LDA model: In our approach, we combine two LDA models [11], [9] to represent users as well as items by mixture of latent topics over time. Topics Over Time (TOT) model [11] is a variant of LDA which models timestamps and the tags in the timestamped items. LDA for collaborative filtering [9] is also a variant of LDA which models users and items over mixtures of latent topics. Our LDA model generates topic-tag distributions and topic-user distributions over time. In our approach, we combine two LDA model [11], [9] to represent users as well as items using a mixture of latent topics over time. The topics over time (TOT) model [11] is a variant of the LDA model that models timestamps and the tags in the timestamped items. The LDA model for collaborative filtering [9] is a variant of the LDA model that models users and items using mixtures of latent topics. Our LDA model generates topic-tag distributions and topic-user distributions over time as follows:

1. For each topic $z \in T$:
   (a) Draw V dimensional multinomials ϕ from a Dirichlet prior β;
   (b) Draw U dimensional multinomials ε from a Dirichlet prior γ;

2. For each item $d \in D$, draw a T dimensional multinomial θ from a Dirichlet prior α;

3. For each tag w_{di} in item d:
   (a) Draw a topic z_{di} from multinomial θ_d;
   (b) Draw a tag w_{di} from multinomial $\phi_{z_{di}}$;
   (c) Draw a timestamp t_{di} from Beta $\psi_{z_{di}}$;

4. For each user u_{di} in item d:
   (a) Draw a topic z_{di} from multinomial θ_d;
   (b) Draw a user u_{di} from multinomial ε_d;

The graphical model for this process is shown in Figure 1, and the symbols in Figure 1 are described in Table I. In Figure 1, only the shaded circles (u, w, and t) are observed data. As shown in the process, the posterior distribution of topics are evaluated by tag, user, and time. The parameterizations are similar to the following:

$$\theta_d|\alpha \sim Dirichlet(\alpha)$$
$$\phi_z|\beta \sim Dirichlet(\beta)$$
$$\varepsilon_z|\gamma \sim Dirichlet(\gamma)$$
$$z_{di}|\theta_d \sim Multinomial(\theta_d)$$
$$w_{di}|\phi_{z_{di}} \sim Multinomial(\phi_{z_{di}})$$

TABLE I
SYMBOLS FOR LDA

Symbol	Description
T	number of topics
D	number of items
V	number of unique tags
U	number of users
N_d	number of tags in item d
α	the hyperparameter of the Dirichlet prior for multinomial θ
β	the hyperparameter of the Dirichlet prior for multinomial ϕ
γ	the hyperparameter of the Dirichlet prior for multinomial ε
ψ	the hyperparameter of Beta
θ_d	the multinomial distribution of tags in the item d
ψ_z	the beta distribution of time specific to topic z
ε_z	the multinomial distribution of users specific to topic z
z_{di}	the topic associated with the ith tag in the item d
w_{di}	the ith tag in item d
u_{di}	the user associated with the ith tag in the item d
t_{di}	the timestamp associated with the ith tag in the item d

$$u_{di}|\varepsilon_{z_{di}} \sim Multinomial(\varepsilon_{z_{di}})$$
$$t_{di}|\psi_{z_{di}} \sim Beta(\psi_{z_{di}})$$

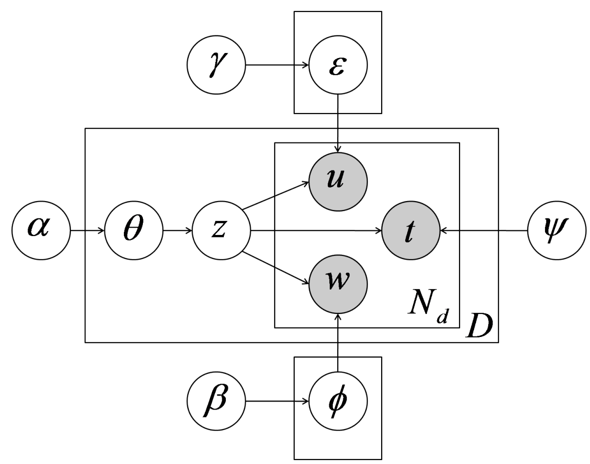

Fig. 1. LDA model for item and tag modeling over time

We now must infer unknown parameters, such as ϕ, ε, and ψ. Because obtaining the exact inference for these parameters is not possible [3], we used Gibbs sampling for obtaining approximate inferences [12], [13], [14]. Using Gibbs sampling, we evaluated the posterior distributions of z and used them to infer ϕ, ε, and ψ. The topic assignment z of a randomly chosen user u, tag w, and time t is sampled from all latent topics according to Equation (1):

$$\begin{aligned} P(z_{di} = k|z_{-di}, R^u, R^w, R^t) \propto \\ \frac{n^t_{R,-i} + \alpha_t}{(\sum_{t=1}^{T} n^t_R + \alpha_t) - 1} \times \frac{n^v_{k,-i} + \beta_v}{\sum_{v=1}^{V} n^v_{k,-i} + \beta_v} \\ \times \frac{n^u_{k,-i} + \gamma_u}{\sum_{u=1}^{U} n^u_{k,-i} + \gamma_u} \times \frac{(1 - t_{di})^{\psi_{k1}^{-1}} t_{di}^{\psi_{k2}^{-1}}}{B(\psi_{k1}, \psi_{k2})} \end{aligned} \quad (1)$$

,where n^u_k is the number of times topic k is assigned for the user u, n_k^w the number of times topic k is assigned to the tag w, and n_R^k the number of times topic k appears in item R; $n_{,-i}$ indicates that the current ith allocation is not counted; B denotes a beta distribution; t_k and s_k^2 in $\psi_{k1} = t_k(\frac{t_k(1-t_k)}{s_k^2} - 1)$ and $\psi_{k2} = 1 - \psi_{k1}$ denote the sample mean and the sample variance of the timestamps belonging to the topic k; and ψ_{k1} is the topic distributions over time. From Equation (1), we can estimate the topic-tag distribution ϕ and the topic-user distribution ϵ using the following formula:

$$\phi_{t,v} = \frac{n_{k,-i}^v + \beta_v}{\sum_{v=1}^{V} n_{k,-i}^v + \beta_v} \tag{2}$$

$$\epsilon_{t,u} = \frac{n_{k,-i}^u + \gamma_u}{\sum_{u=1}^{U} n_{k,-i}^u + \gamma_u} \tag{3}$$

As topic models are typically sensitive to hyperparameters, it is important to obtain the correct values for the hyperparameters. After finding that the sensitivity to hyperparameters was not very strong in our model, we used fixed symmetric Dirichlet distributions ($\alpha = 50/T$, $\beta = 0.1$, and $\gamma = 0.1$).

C. Time-based Similarity Weight Calculation

The previous section described the manner in which we obtained topic-user and topic-tag distributions. By obtaining topic-user distributions, we can understand which users have similar topic distributions and perform user grouping by using clustering methods, such as k-means clustering, by which each user is an entity whose attributes are topics. As described in the introduction, when a recommendation system suggests a new item to a specific user at a specific point in time, the system identifies which of the user's current items is most similar to the new item and, if their degree of similarity reaches a threshold, recommends the new item to the user. As user interests can change over time, our model may contain users who have different tag distributions over time and, consequently, different topic distributions over time. Thus, given topic distributions at an arbitrary point in time, certain topic distributions will have a higher level of similarity with the distributions than will topic distributions at other points in time. Taking into account different topic distributions according to time, we can calculate the similarity among topic distributions according to time periods and adopt the similarity as a weight in order to suggest a new item to a user.

Users in the same group who have similar interest dynamics over time may be grouped together. To do so, we created user groups based on their topic distributions and calculated similarity weights over time for each group, as we had in the previous section. After identifying the user groups, we collected user items associated with each group and then divided the collected items according to the time period (in this case one month). Defining each item as a topic vector, we identified the tags associated with each item and the topic-tag distributions for all the tags. Converting new items into topic vectors, we incorporated topic vectors from the new items and topic vectors of the group from each month into the dataset in order to measure the similarity of the weights. For example, if one group's items existed over 12 months, each group had 12 datasets. From the dataset, we measured the topic similarity over time using Equation (4). As a result, we obtained several topic similarity values over time for each group that we termed *group similarity weights over time*.

$$weight(g,t) = \frac{\overrightarrow{x_{target}} \cdot \overrightarrow{x_{g,t}}}{\|\overrightarrow{x_{target}}\| \|\overrightarrow{x_{g,t}}\|} \tag{4}$$

, where $\overrightarrow{x_{target}}$ denotes the topic vector in the target time $target$, $\overrightarrow{x_{g,t}}$ denotes the topic vector in time $t \in TS$ for group g, and TS is the set of time slots of the data set.

D. Recommendation System

In this section, we explain two recommendation approaches. When a recommendation system suggests a new item to a specific user, the system finds user's item which the similarity with the new item is the highest. There are, however, differences in using temporal information for two approaches. The first approach is a static recommendation: it does not employ temporal information. The second approach is a dynamic recommendation.

1) Static topic based recommendation: The topic-based recommendation system described in this section served as the basis for our approach. First, we converted each item into a topic vector, as described in the previous section. To determine whether to recommend a new item to a user, we identified which of his or her items has the greatest level of similarity with the new item and, if the level of similarity reached a threshold, recommended the new item to the user. To determine the level of similarity, we determined the cosine similarity between items using Equation (5):

$$sim(\overrightarrow{y_i}, \overrightarrow{y_j}) = \frac{\overrightarrow{y_i} \cdot \overrightarrow{y_j}}{\|\overrightarrow{y_i}\| \|\overrightarrow{y_j}\|} \tag{5}$$

, where $\overrightarrow{y_i}$ and $\overrightarrow{y_j}$ are topic vectors for item i and j.

2) Dynamic topic based recommendation with group similarity weight over time: To recommend a new item to a user using group similarity weights over time from the previous section, we first identified the group with which the user was initially associated. We then calculated the similarity between the new item and the user's items employing group similarity weights over time, which differ according to the user's group and the group item's time period. By using group similarity weights over time in Equation (6) to calculate the similarity between items, we could identify which item had the highest level of similarity with the new item and, if the level reached a threshold value, recommend the item.

$$sim(\overrightarrow{y_i}, \overrightarrow{y_{j,g,t}}) = weight(g,t) * \frac{\overrightarrow{y_i} \cdot \overrightarrow{x_{j,g,t}}}{\|\overrightarrow{y_i}\| \|\overrightarrow{y_{j,g,t}}\|} \tag{6}$$

, where $\overrightarrow{y_i}$ is a topic vector for item i, $\overrightarrow{y_{j,g,t}}$ is a topic vector for item j in the time slot t and in the group g where the user is associated and $weight(g,t)$ is a similarity weight for time t in the group g.

IV. Experimental Evaluation

A. Data Retrieval and Preprocessing

We need to show that our dynamic item recommendation is more effective than static item recommendation. To do this, we have experimentally retrieved real data from social networks. Our dataset needs to have three entities: user, item, and tag. In our dataset, users have collections of items which are selected by them or recommended to them in order to show their preferences explicitly and items are described by tags. Also, we need to track the selection of the items over time in order to manipulate temporal information.

To conduct our experiments, we used datasets from Flickr, a photo-sharing network that allows users to upload photos, describe them with tags, and set other users' photos as their favorite photos. Considering that setting a photo as a favorite is a strong sign of user preference, our objective in this experiment was to recommend photos set as favorite photos by users given their previous favorite photo datasets. The data that we collected from Flickr consisted of users' information, users' favorite photos, and the tags associated with the favorite photos. In our experiments, all the photos were favorite photos and considered items in our approach. Our dataset contained 5,821 users, 1,183,398 photos and 13,721,075 favorite tags regarding photos chosen by users as their favorite photos between June 1, 2008 and June 30, 2009.

We sorted based on tf-idf weights to prune the data to make it more manageable and to suppress noise. For each user, we pruned tags that fell below top 0.5% by tf-idf weights, and retained only those photos that had at least one of the top 0.5% tags for each user. After pre-processing the data, we executed Gibbs sampling for the LDA model. The topic number for the LDA model sampling was 50. As a result, we retrieved topic-user and topic-tag distributions over time. Using user-topic distributions, we converted each user into a topic vector that we used to create user groups by the clustering algorithm like k-means clustering, with the value of k set at 20. After creating user groups, we identified user items associated with the group by each month (having set the time period to one month) to calculate the similarity weight for each group and each time period.

We then created a training set that included topic vectors of new items and user items within a specific month. From the training set, we determined a similarity weight for the specific time period of the group. To determine the similarity weight, we analyzed 500 user photos chosen as favorites between June 2008 and May 2009, and considered the photos chosen in June 2009 to be the current favorite photos. When we applied our recommendation system to 3,821 users to test our system using the dataset, we found that the ratio of positive photos (labeled as recommended) and negative photos (labeled as non-recommended) of the 186,408 photos that we evaluated was 1:5.

B. Precision Evaluation of Static and Dynamic Systems

As the item space in social networks is almost infinite, it is impossible to retrieve all possible items in which users may be interested. However, it is possible and valuable to recommend items that exactly match users' interests by considering the precision results of the approaches in which recall is greater than 0.1. Regarding which of the approaches to recommend based on the results obtained in the experiments and the precision rate, we first recommend photos without using any weights, then recommend photos using group similarity weights.

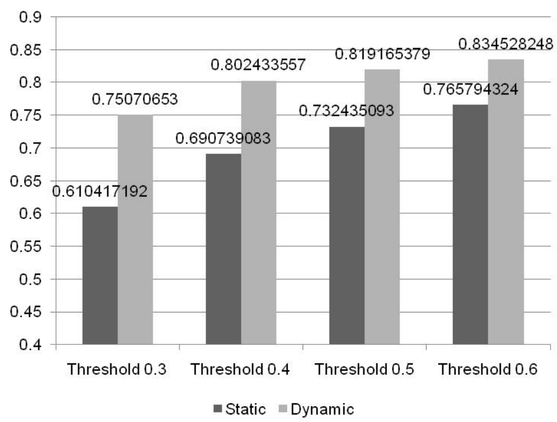

Fig. 2. Precision Rates over Different Settings

In Figure 2, the precision results of two recommendation approaches are displayed. The X-axis denotes the threshold for the similarity between the new photo and the user's photo. For example, if the threshold is 0.3, the similarity between the new photo and the most similar user's previous photo is greater than 0.3, so the new photo is recommended to the user. The Y-axis denotes the precision rate for the recommendation. Each column denotes a recommendation approach. The first column denotes the static recommendation and the second column denotes the dynamic recommendation. The precision rates of the static recommendation are 0.61, 0.69, 0.73, and 0.76 given the thresholds of 0.3, 0.4, 0.5, and 0.6 respectively. The precision rates of the recommendation using group weights are 0.75, 0.80, 0.81, and 0.83 given thresholds. The dynamic recommendation system outperforms the static recommendation system, confirming our hypothesis that consideration of the phenomenon that user interests change over time is necessary to increase the precision with which recommendations are made by a system.

C. Top-K Precision Evaluation of Static and Dynamic Systems

In this section, we adopt another evaluation rate: top-k precision rate. In the recommended data, we sort recommended images by their similarities for each user. Then we collect sorted recommended images for each user. Then we pick top-k images and calculate precision rates from top-k images for each user. In real system, users do not want a huge number of recommended data, they just want a small number of recommended data. In this assumption, the adoption of top-k precision rate is reasonable.

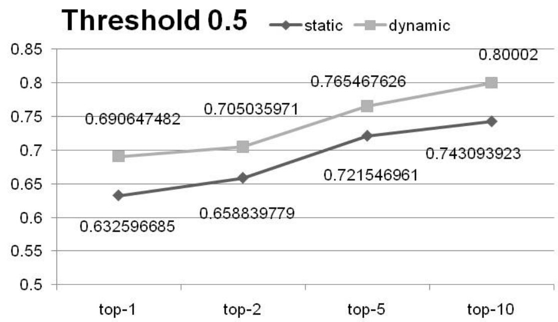

Fig. 3. Similarity Weights over Time

As shown in Figure 3, the top-k precision results of two recommendation approaches are displayed. The X-axis denotes the top-k number. Each slot denotes top-1, top-2, top-5, and top-10 respectively. The Y-axis denotes the precision rate. The data is recommended with a similarity threshold 0.5. The top-k precision rates for dynamic recommendation are 0.69, 0.70, 0.76, and 0.80 for each top-k while the top-k precision rates for static recommendation are 0.62, 0.65, 0.72, and 0.74 respectively. This graph also confirms that the dynamic recommendation outperforms the static recommendation.

V. Conclusion

The recommendation systems used within social networks must address the phenomenon that user interests change over time. We addressed this phenomenon by developing and testing a recommendation system that matches user and group interests over time by topics extracted by tags associated with items to make recommendations. The data in our approach consists of tuples of a user, a set of favorite items, and associated tags. As manipulating a dataset is computationally difficult due to its inherent noisiness and large feature space, we pre-processed data by calculating tf-idf weights for each tag of each user and, after sorting the tags by tf-idf weight, retained only those tags with high weights for making recommendations. Although this preprocessing reduced the number of tags, too many tags remained to make recommendations. To further reduce the complexity and model items by topics over time, we applied the LDA model to identify latent topics from items and tags, which allowed us to derive determine topic-user and topic-tag distributions over time. After grouping the users by topic-user distributions and evaluating the similarity weights given the groups and time, we employed the similarity weights to calculate the level of similarity between users' items and new items for making recommendations. Our results proved promising. When we compared the precision rates on test data with the different systems: a static system not using weights; and a dynamic system using group similarity weights over time: We found that the system using the group similarity weights over time helped improve the precision rate.

In conclusion, we proposed an approach for item recommendation by examining tag vocabulary over time. Our results demonstrated that tags can serve as useful keys in identifying user preferences and that gaining understanding of trends in user interests over time is essential for better recommendation results. We plan to extend our work in the future by adding more temporal aspects of item recommendation, such as identifying periodic or seasonal trends and extending our focus from discovering monthly trends to discovering daily trends, to provide more up-to-date recommendations for users. Also, we plan to apply temporal information to user group generation in order to grasp the actual shape of the group and use the information for recommendation.

References

[1] U. Hanani, B. Shapira, and P. Shoval, "Information Filtering: Overview of Issues, Research and Systems," *User Modeling and User-Adapted Interaction*, vol. 11, pp. 203–259, 2001.

[2] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using collaborative filtering to weave an information tapestry," *Commun. ACM*, vol. 35, no. 12, pp. 61–70, 1992.

[3] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, 2003.

[4] Y. Zhen, W.-J. Li, and D.-Y. Yeung, "Tagicofi: Tag informed collaborative filtering," in *Proceedings of the 3rd ACM Conference on Recommender Systems (RecSys '09)*, New York City, New York, USA, Oct. 22–25 2009, pp. 69–76.

[5] S. Sen, J. Vig, and J. Riedl, "Tagommenders: connecting users to items through tags," in *WWW '09: Proceedings of the 18th international conference on World wide web*, 2009, pp. 671–680.

[6] M. de Gemmis, P. Lops, G. Semeraro, and P. Basile, "Integrating tags in a semantic content-based recommender," in *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, 2008, pp. 163–170.

[7] Z. Guan, C. Wang, J. Bu, C. Chen, K. Yang, D. Cai, and X. He, "Document recommendation in social tagging services," in *WWW '10: Proceedings of the 19th international conference on World wide web*, 2010, pp. 391–400.

[8] S. Siersdorfer and S. Sizov, "Social recommender systems for web 2.0 folksonomies," in *HT '09: Proceedings of the 20th ACM conference on Hypertext and hypermedia*, 2009, pp. 261–270.

[9] Y. Guo and J. B. Joshi, "Topic-based personalized recommendation for collaborative tagging system," in *HT '10: Proceedings of the 21st ACM conference on Hypertext and hypermedia*, 2010, pp. 61–66.

[10] K. S. Jones, "A statistical interpretation of term specificity and its application in retrieval," *Journal of Documentation*, vol. 28, pp. 11–21, 1972.

[11] X. Wang and A. McCallum, "Topics over time: a non-markov continuous-time model of topical trends," in *KDD '06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006, pp. 424–433.

[12] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. Suppl 1, pp. 5228–5235, April 2004. [Online]. Available: http://dx.doi.org/10.1073/pnas.0307752101

[13] M. Steyvers, P. Smyth, M. Rosen-Zvi, and T. Griffiths, "Probabilistic author-topic models for information discovery," in *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2004, pp. 306–315.

[14] X. Wei and W. B. Croft, "Lda-based document models for ad-hoc retrieval," in *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, 2006, pp. 178–185.