

Learn to Personalized Image Search From the Photo Sharing Websites

Jitao Sang, Changsheng Xu, *Senior Member, IEEE*, and Dongyuan Lu

Abstract—Increasingly developed social sharing websites like Flickr and Youtube allow users to create, share, annotate, and comment medias. The large-scale user-generated metadata not only facilitate users in sharing and organizing multimedia content, but provide useful information to improve media retrieval and management. Personalized search serves as one of such examples where the web search experience is improved by generating the returned list according to the modified user search intents. In this paper, we exploit the social annotations and propose a novel framework simultaneously considering the user and query relevance to learn to personalized image search. The basic premise is to embed the user preference and query-related search intent into user-specific topic spaces. Since the users' original annotation is too sparse for topic modeling, we need to enrich users' annotation pool before user-specific topic spaces construction. The proposed framework contains two components: 1) a ranking-based multicorrelation tensor factorization model is proposed to perform annotation prediction, which is considered as users' potential annotations for the images; 2) we introduce user-specific topic modeling to map the query relevance and user preference into the same user-specific topic space. For performance evaluation, two resources involved with users' social activities are employed. Experiments on a large-scale Flickr dataset demonstrate the effectiveness of the proposed method.

Index Terms—Personalized image search, social annotation, tensor factorization, topic model.

I. INTRODUCTION

KEYWORD-BASED search has been the most popular search paradigm in today's search market. Despite simplicity and efficiency, the performance of keyword-based search is far from satisfying. Investigation has indicated its poor user experience—on Google search, for 52% of 20 000 queries, searchers did not find any relevant results [1]. This is due to two reasons: 1) queries are in general short and nonspecific, e.g., the query of “IR” has the interpretation of both information retrieval and infra-red, and 2) users may have different intentions for the same query, e.g., searching for “jaguar” by a car fan has a completely different meaning

Manuscript received June 21, 2011; revised October 28, 2011; accepted November 09, 2011. Date of publication December 22, 2011; date of current version July 13, 2012. This work was supported in part by the National Program on Key Basic Research Project (973 Program, Project 2012CB316304) and by the National Natural Science Foundation of China under Grant 90920303 and Grant 61003161. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jinhui Tang.

J. Sang and C. Xu are with the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China, and also with the China-Singapore Institute of Digital Media, 119615 Singapore (e-mail: jtsang@nlpr.ia.ac.cn; csxu@nlpr.ia.ac.cn).

D. Lu is with the State Key Laboratory of Intelligent Control and Management of Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China (e-mail: dongyuan.lu@ia.ac.cn).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMM.2011.2181344

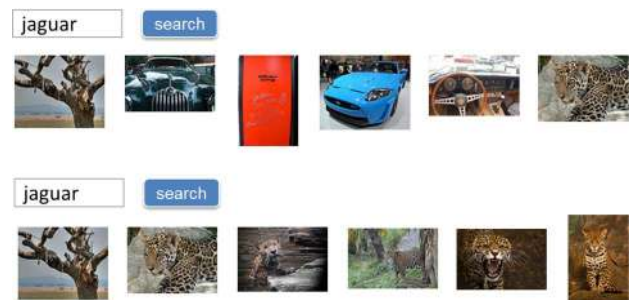


Fig. 1. Toy example for (top) non-personalized and (bottom) personalized search results for the query “jaguar”.

from searching by an animal specialist. One solution to address these problems is *personalized search*, where user-specific information is considered to distinguish the exact intentions of the user queries and rerank the list results. Given the large and growing importance of search engines, personalized search has the potential to significantly improve searching experience.

Compared with non-personalized search, in personalized search, the rank of a document (web page, image, video, etc.) in the result list is decided not only by the query, but by the preference of user. Fig. 1 shows a toy example for non-personalized and personalized image search results. The non-personalized search returned results only based on the query relevance and displays jaguar car images as well as wild cat on the top. While personalized search consider both query relevance and user preference, therefore the personalized results from an animal lover rank the leopard images on the top. This provides a natural two-step solution scheme. Most of the existing work [2]–[5] follow this scheme and decompose personalized search into two steps: computing the non-personalized relevance score between the query and the document, and computing the personalized score by estimating the user's preference over the document. After that, a merge operation is conducted to generate a final ranked list. While this two-step scheme is extensively utilized, it suffers from two problems. 1) The interpretation is less straight and not so convinced. The intuition of personalized search is to rank the returned documents by estimating the user's preference over documents under certain queries. Instead of directly analyzing the user-query-document correlation, the existing scheme approximates it by separately computing a query-document relevance score and a user-document relevance score. 2) How to determine the merge strategy is not trivial.¹ In this paper, we simultaneously considers the

¹Typically a weighting parameter will be optimized to balance the two scores [2], or the learnt user preference is used to rerank the query relevance-based original list [5].

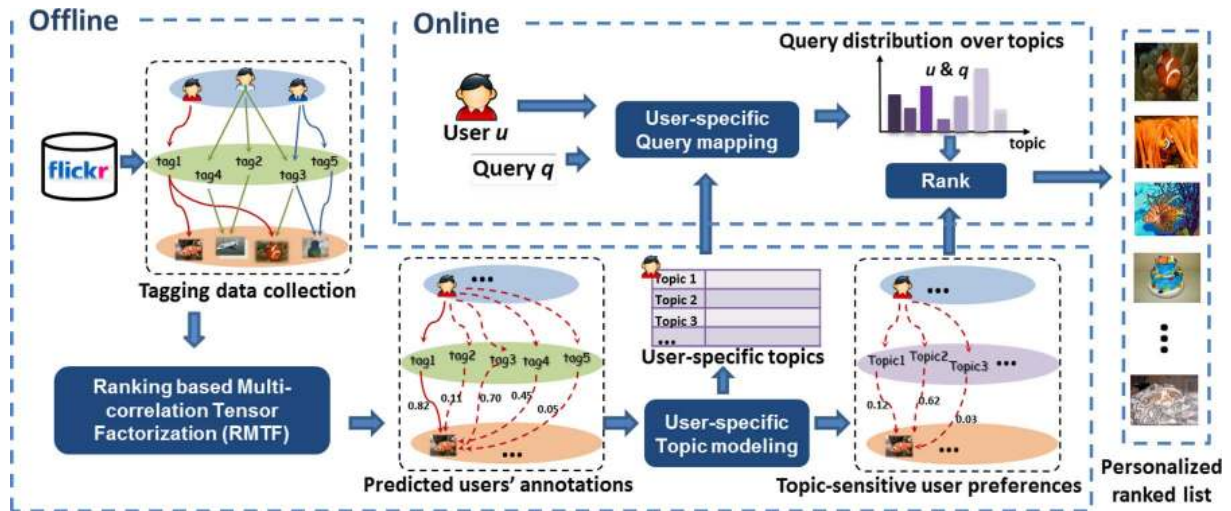


Fig. 2. Proposed framework.

user and query dependence and present a novel framework to tackle the personalized image search problem.

To investigate on user preference and perform user modeling, the popular social activity of tagging is considered. Collaborative tagging has become an increasingly popular means for sharing and organizing resources, leading to a huge amount of user-generated annotations. Online photo sharing websites, such as Flickr, Picasa, Zoomr, and Pinterest allow users as owners, taggers, or commenters for their contributed contents to interact and collaborate with each other in a social media dialogue. Various researchers have investigated the applicability of social annotations to improve web search [6]–[10]. Recently, social annotations are employed for automatic evaluation of personalized search [2], [11], [3]. A fundamental assumption is that, *the users' tagging actions reflect their personal relevance judgement*. For example, if a user tagged “festival” to an image, it is probable that the user will consider this image as relevant if he/she issues “festival” as a query. Illustrated by this, the intuition of this paper is that if the users' annotations to the images are available, we can directly estimate the users' preference under certain queries. The fact is that the original annotations available is not enough for user preference mining. Therefore, we transfer the problem of personalized image search to users' annotation prediction. Moreover, as queries and tags do not follow simple one-to-one relationship, we build user-specific topic spaces to exploit the relations between queries and tags.

A. Framework

The framework of this paper is shown in Fig. 2. It contains two stages: offline model training stage and online personalized search response stage.

For the offline stage, three types of data including users,² images and tags as well as their ternary interrelations and intra-relations are first collected.³ We then perform users'

²We use Flickr, the popular photo sharing website, for our experiments. The information of users, images and tags can be acquired from the Flickr API: <http://www.flickr.com/services/api>.

³We show a running example consisting of three users, five tags and four images.

annotation prediction. Many methods [12]–[14] for tag recommendation and prediction have been proposed in social bookmark sites, e.g., Bibsonomy, Del.icio.us, Last.fm, etc. Since the photo sharing websites utilize a different tagging mechanism that repetitive tags are not allowed for unique images, besides the common noisy problem, it has more severe sparsity problem than other social tagging systems.⁴ To alleviate the sparsity and noisy problem, we present a novel method named *ranking-based multicorrelation tensor factorization* (RMTF) to better leverage the observed tagging data for users' annotation prediction. Zhu *et al.* [15] has demonstrated that the semantic space spanned by image tags can be approximated by a smaller subset of salient words from the original space. Illustrated by this, we employ low-rank approximation to extract the compact representation for image, tag and user, and at the same time reconstruct the user-image-tag ternary relations for annotation prediction. With the observed user-tag-image ternary relations as input, the reconstructed ternary relations can be viewed as users' potential annotations for the images.

Following the assumption we mentioned in the introduction, we can straightly utilize the predicted user annotations for personalized image search, i.e., if a user has a high probability to assign the tag t to an image, the image should be ranked higher when the user issues query t . However, this formulation has two problems: 1) it is unreasonable to assign the query to a single tag in the tag vocabulary, e.g., when a user searches “cheerdance”, he/she would like the images that he/she annotated with semantic related tag “cheerleader” are also ranked higher, and 2) there are variations in individual user's tagging patterns and vocabularies, e.g., the tag “jaguar” from an animal specialist should be related to “leopard”, while a car fan will consider “jaguar” more related to “autos”. To address the two problems, we perform *user-specific topic modeling* to build the semantic topics for each user. The user's annotation for an image is viewed as *document*. The individual tag to the image is *word*. User's annotations for all the images constitute the *corpus*. As the original annotation is too sparse for topic mod-

⁴We conduct the sparsity investigation between Flickr, the publication tagging website Bibsonomy and the music sharing websites Last.fm in Section V.

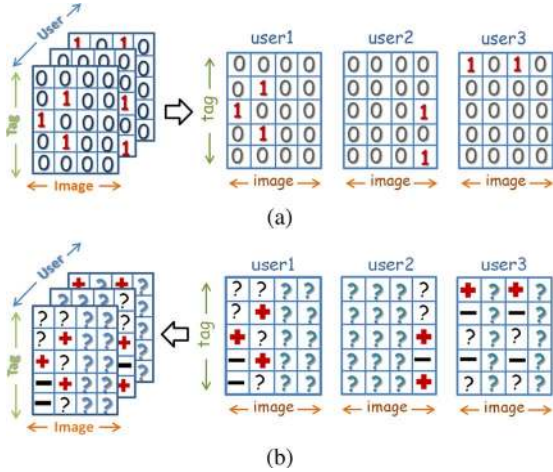


Fig. 3. Tagging data interpretation. (a) 0/1 scheme. (b) Ranking scheme.

eling, we use the reconstructed ternary relations as the document collections. The user’s topic distribution per image can be considered as his/her preference over the image on the learned user-specific topic space. Therefore, after the offline stage, two outcomes are stored in the system, the user-specific topics and topic-sensitive user preferences.

For the online stage, when a user u submits a query q , we first map the query q to user u -specific topics. The query distribution is then sent to the rank module and employed as the weight on topics to calculate the user u ’s topic-sensitive preferences over the images. Finally, the images are ranked according to the calculated user’s preferences, which simultaneously considers the query and user information.

The contributions of this paper are summarized as three-folds.

- We propose a novel personalized image search framework by simultaneously considering user and query information. The user’s preferences over images under certain query are estimated by how probable he/she assigns the query-related tags to the images.
- A ranking based tensor factorization model named RMTF is proposed to predict users’ annotations to the images.
- To better represent the query-tag relationship, we build user-specific topics and map the queries as well as the users’ preferences onto the learned topic spaces.

II. RELATED WORK

In recent years, extensive efforts have been focusing on personalized search. Regarding the resources they leveraged, explicit user profile [17], relevance feedback [18], user history data (browsing log [19], click-through data [20], [21], and social annotations [11], [8], [4], etc.), context information [23] (time, location, etc.) and social network [1], [3], [16] are exploited. For the implementation there are two primary strategies [24], query refinement and result processing. In the following we review the related work by the strategy they used.

Query refinement, also called query expansion, refers to the modification to the original query according to the user information. It includes augmenting the query by other terms [18], [25] and changing the original weight of each query term [26].

Kraft *et al.* [18] utilized the search context information collected from users’ explicit feedback to enrich the query terms. Chirita *et al.* [25] proposed five generic techniques for providing expansion terms, ranging from term and expression level analysis up to global co-occurrence statistics and external thesauri. While, Teevan *et al.* [26] reassigned the weights of original query terms using BM25 weighting scheme to incorporate user interests as collected by their desktop indexes. We do not explicitly perform query refinement in this paper. However, mapping the queries into user-specific topic spaces can be considered as implicit query refinement.

Result processing can be further classified into result filtering and reranking. Result filtering aims to filter irrelevant results that are not of interest to a particular user [27]. While, result reranking focuses on reordering the results by the degree of users’ preferences estimated. Since our work falls into this category, we mainly review the related work on result reranking. Chirita *et al.* [17] conducted an early work by reranking the search results according to the cosine distance between each URL and user interest profiles constructed. Qiu *et al.* [21] extended Topic-Sensitive PageRank by incorporating users’ preference vectors. By aggregating the search results from multiple search engines, Liu *et al.* [22] introduced a new method for visual search reranking called CrowdReranking. A typical work is performed by Xu *et al.* [2], in which the overall ranking score is not only based on term similarity matching between the query and the documents but also topic similarity matching between the user’s interests and the documents’ topics. In the similar spirits, Cai [4] formalized query and user relevance measurement separately as fuzzy requirement satisfaction problem. Lu *et al.* [5] utilized a coclustering method to extract latent interest dimensions, and rerank the images by combining latent interest based user preference and query relevance. In our work, there is also a topic space to model user preference. However, regarding the variations in user’s tagging vocabularies, we build user-specific topics and calculate topic-sensitive user preference over images, which differentiate our work from Xu [2] and Lu’s [5]. Besides, most of the existing work decompose the overall ranking score into query relevance and user preference and generate two separated ranked list. While in this paper, we map the queries into the same user-specific topic space and directly compute the users’ preference under certain queries.

III. RANKING BASED MULTICORRELATION TENSOR FACTORIZATION

In this section, we present the algorithm for annotation prediction. Table I lists the key notations used in this paper. There are three types of entities in the photo sharing websites. The tagging data can be viewed as a set of triplets. Let $\mathcal{U}, \mathcal{I}, \mathcal{T}$ denote the sets of users, images, tags and the set of observed tagging data is denoted by $\mathcal{O} \subset \mathcal{U} \times \mathcal{I} \times \mathcal{T}$, i.e., each triplet $(u, i, t) \in \mathcal{O}$ means that user u has annotated image i with tag t . The ternary interrelations can then constitute a three dimensional tensor $\mathcal{Y} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}| \times |\mathcal{T}|}$, which is defined as

$$y_{u,i,t} = \begin{cases} 1, & \text{if } (u, i, t) \in \mathcal{O} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

TABLE I
LIST OF KEY NOTATIONS

Symbol	Description
\mathcal{Y}, \mathcal{C}	user-image-tag tensor, and core tensor
U, I, T	represent user, image, tag factor matrices
$\mathbb{U}, \mathbb{I}, \mathbb{T}$	sets of users, images and tags, respectively
u, i, t	represent user, image, tag index
$\mathbf{u}, \mathbf{i}, \mathbf{t}$	represent user, image, tag feature vectors

Fig. 3(a) shows the tensor constructed from the running example in Fig. 2.

Predicting the users' annotations to the images are related to reconstructing the user-tag-image ternary interrelations. We use Tucker decomposition [31], a general tensor factorization model, to perform the low-rank approximation. In Tucker decomposition, the tagging data \mathcal{Y} are estimated by three low rank matrices and one core tensor

$$\hat{\mathcal{Y}} := \mathcal{C} \times_u U \times_i I \times_t T \quad (2)$$

where \times_n is the tensor product of multiplying a matrix on mode n . Each matrix ($U \in \mathbb{R}^{|\mathbb{U}| \times r_U}$, $I \in \mathbb{R}^{|\mathbb{I}| \times r_I}$, $T \in \mathbb{R}^{|\mathbb{T}| \times r_T}$) corresponds to one factor. The core tensor $\mathcal{C} \in \mathbb{R}^{r_U \times r_I \times r_T}$ contains the interactions between the different factors. The ranks of decomposed factors are denoted by r_U, r_I, r_T and this is called *rank*-(r_U, r_I, r_T) Tucker decomposition. Under Tucker decomposition, we need to design appropriate objective function to optimize the latent factors U, I, T, C and then calculate the reconstructed tensor by (2).

In this paper, a model named RMTF is proposed to design the objective function. To better leverage the observed tagging data, we first introduce a novel ranking based optimization scheme for representation of the tagging data. Then the multiple intra-relations among users, images and tags are utilized as the smoothness constraints to tackle the sparsity problem.

A. Ranking Based Optimization Scheme

A direct way to approximate \mathcal{Y} is to minimize the sum of point-wise loss on \mathcal{Y}

$$\min_{U, I, T, C} \sum_{(\tilde{u}, \tilde{i}, \tilde{t}) \in |\mathbb{U}| \times |\mathbb{I}| \times |\mathbb{T}|} (\hat{y}_{\tilde{u}, \tilde{i}, \tilde{t}} - y_{\tilde{u}, \tilde{i}, \tilde{t}})^2 \quad (3)$$

where $\hat{y}_{\tilde{u}, \tilde{i}, \tilde{t}} = \mathcal{C} \times_u \mathbf{u}_{\tilde{u}} \times_i \mathbf{i}_{\tilde{i}} \times_t \mathbf{t}_{\tilde{t}}$. As this optimization scheme tries to fit to the numerical values of 1 and 0, we refer it as the *0/1 scheme*.

However, under the situation of social image tagging data, the semantics of encoding all the unobserved data as 0 are incorrect, which is illustrated with the running example.

- Firstly, the fact that *user3* has not given any tag to *image2* and *image4* does not mean *user3* considering all the tags are bad for describing the images. Maybe he/she does not want to tag the image or has no chance to see the image.⁵
- Secondly, *user1* annotates *image1* with only *tag3*. It is also unreasonable to assume that other tags should not be annotated to the image, as some concepts may be missing in

⁵We call the triplets like $(u_3, i_2, :)$ and $(u_3, i_4, :)$ as neutral triplets.

the user-generated tags and individual user may not be familiar to all the relevant tags in the large tag vocabulary.

According to the optimization function in (3), *0/1 scheme* tries to predict 0 for both cases. To address the above two issues, in this paper, we present a ranking optimization scheme which intuitively takes the user tagging behaviors into consideration.

Firstly, we note that only the qualitative difference is important and fitting to the numerical values of 1 and 0 is unnecessary. Therefore, instead of solving an point-wise classification task, we formulate it as a ranking problem which uses tag pairs within each user-image combination (u, i) as the training data and optimizes for correct ranking.

We provide some notations for easy explanation. Each user-image combination (u, i) is defined as a *post*. The set of observed posts is denoted as $\mathbb{P}_\mathbb{O}$

$$\mathbb{P}_\mathbb{O} = \{(u, i) | \exists t \in \mathbb{T}, y_{u, i, t} = 1\}. \quad (4)$$

Note that the ranking optimization is performed over each post and within each post (u, i) a positive tag set $\mathbb{T}_{u, i}^+$ and a negative tag set $\mathbb{T}_{u, i}^-$ are desired to construct the training pairs. We assume that any tag $t \in \mathbb{T}_{u, i}^+$ is a better description for image i than all the tags $t \in \mathbb{T}_{u, i}^-$. The pairwise ranking relationships can be denoted as

$$\hat{y}_{u, i, t_1} > \hat{y}_{u, i, t_2} \Leftrightarrow t_1 \in \mathbb{T}_{u, i}^+ \wedge t_2 \in \mathbb{T}_{u, i}^-. \quad (5)$$

The optimization criterion is to minimize the violation of the pairwise ranking relationships in the reconstructed tensor $\hat{\mathcal{Y}}$, which leads to the following objective:

$$\min_{U, I, T, C} \sum_{(\tilde{u}, \tilde{i}) \in \mathbb{P}_\mathbb{O}} \left(\sum_{t^+ \in \mathbb{T}_{\tilde{u}, \tilde{i}}^+} \sum_{t^- \in \mathbb{T}_{\tilde{u}, \tilde{i}}^-} f(\hat{y}_{\tilde{u}, \tilde{i}, t^-} - \hat{y}_{\tilde{u}, \tilde{i}, t^+}) \right) \quad (6)$$

where $f: \mathbb{R} \rightarrow [0, 1]$ is a monotonic increasing function (e.g., the logistic sigmoid function or Heaviside function).

Secondly, for the training pair determination. The neutral triplets constitute a set \mathbb{M}

$$\mathbb{M} = \{(u, i, t) | (u, i) \notin \mathbb{P}_\mathbb{O}\}. \quad (7)$$

It is arbitrary to treat the neutral triplets as either positive or negative and we remove all the triplets in \mathbb{M} from the learning process [filled by bold question marks in Fig. 3(b)].

We then consider two characteristics of the user tagging behaviors to choose $\mathbb{T}_{u, i}^+$ and $\mathbb{T}_{u, i}^-$. On one hand, some concepts may be missing in the user-generated tags. We assume that the tags co-occurring frequently are likely to appear in the same image (we call it *context-relevant*). On the other hand, users will not bother to use all the relevant tags to describe the image. The tags *semantic-relevant* with the observed tags are also the potential good descriptions for the image.

To perform the idea, we build a tag affinity graph W^T based on tag semantic and context intra-relations (detailed in Section III-B). The tags with the k -highest affinity values are considered semantic-relevant and context-relevant. Given a




Fig. 4. (a) Cumulative distribution of document ratio w.r.t. the number of taggers for Flickr and Del.icio.us. (b) Number of tagged images per user for Flickr.

post $(u, i) \in \mathbb{P}_{\mathcal{O}}$, the observed tags constitute a positive tag set [the corresponding triplets are filled by plus signs in Fig. 3(b)]

$$\mathbb{T}_{u,i}^+ = \{t | (u, i) \in \mathbb{P}_{\mathcal{O}} \wedge y_{u,i,t} = 1\}. \quad (8)$$

Instead of adding semantic and context-relevant tags into the positive set $\mathbb{T}_{u,i}^+$, we only keep the unobserved tags semantic-**irrelevant** and context-**irrelevant** to any of the observed tags $t \in \mathbb{T}_{u,i}^+$, to form the negative tag set:

$$\mathbb{T}_{u,i}^- = \{t | (u, i) \in \mathbb{P}_{\mathcal{O}} \wedge y_{u,i,t} \neq 1 \wedge t \notin \mathbb{N}_{\mathbb{T}_{u,i}^+}\} \quad (9)$$

where $\mathbb{N}_{\mathbb{T}_{u,i}^+}$ means the set of tags relevant to the annotated tags in post (u, i) . Then $t_4, t_5 \in \mathbb{T}_{u_1, i_1}^-$ when $tag1$ and $tag2$ are relevant to $tag3$. The minus signs indicate the filtered negative triplets in Fig. 3(b). The triplets corresponding to tags $t \in \mathbb{N}_{\mathbb{T}_{u,i}^+}$ are also removed from the learning process and filled by plain question marks.

B. Multicorrelation Smoothness Constraints

Photo sharing websites differentiate from other social tagging systems by its characteristic of self-tagging: most images are only tagged by their owners. Fig. 4(a) shows the #tagger statistics for Flickr and the webpage tagging system Del.icio.us. We can see that in Flickr, 90% images have no more than four taggers and the average number of tagger for each image is about 1.9. However, the average tagger for each webpage in Del.icio.us is 6.1. The severe sparsity problem calls for external resources to enable information propagation.

In addition to the ternary interrelations, we also collect multiple intra-relations among users, images and tags. These intra-relations constitute the affinity graphs $W^U \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{U}|}$, $W^I \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{I}|}$, and $W^T \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$, respectively. We assume that two items with high affinities should be mapped close to each other in the learnt factor subspaces. In the following, we first introduce how to construct the tag affinity graph,⁶ and then incorporate them into the tensor factorization framework.

Tag affinity graph W^T . To serve the ranking based optimization scheme, we build the tag affinity graph based on the tag semantic relevance and context relevance. The context rel-

evance of tag t_m and t_n is simply encoded by their weighted co-occurrence in the image collection

$$t_{m,n}^c = \frac{n(t_m, t_n)}{n(t_m) + n(t_n)}. \quad (10)$$

For tag semantic relevance, we estimate the semantic relevance between tag t_m and t_n based on their WordNet distance

$$t_{m,n}^s = \frac{2 \cdot IC(lcs(t_m, t_n))}{IC(t_m) + IC(t_n)} \quad (11)$$

where $IC(\cdot)$ is the information content of tag, and $lcs(t_i, t_j)$ is their least common subsumer in the WordNet taxonomy. The tag affinity graph is constructed as

$$W_{m,n}^T = \lambda_c t_{m,n}^c + \lambda_s t_{m,n}^s \quad (12)$$

where $\lambda_c + \lambda_s = 1$, λ_c and λ_s are the weights of context relevance and semantic relevance (in the experiment, we choose $\lambda_c = 0.9$ and $\lambda_s = 0.1$).

The affinity graphs can be utilized as the regularization terms to impose smoothness constraints for the latent factors. Take the image affinity graph W^I and the image factor matrix I as example, the regularization term is

$$\sum_{m=1}^{|\mathcal{I}|} \sum_{n=1}^{|\mathcal{I}|} W_{m,n}^I \|\mathbf{i}_m - \mathbf{i}_n\|^2 \quad (13)$$

where $\|\cdot\|^2$ denotes the Frobenius norm. The basic idea is to make the latent representations of two images as close as possible if there exists strong affinity between them. We can achieve this by minimizing the trace of $I^T L_I I$. L_I is the Laplacian matrix for the image affinity matrix W^I . We can build similar regularization terms for the user and tag factors. Combining with (6), we obtain the following overall objective function:

$$\begin{aligned} \min_{U, I, T, C} g = & \sum_{(\tilde{u}, \tilde{i}) \in \mathbb{P}_{\mathcal{O}}} \left(\sum_{t^+ \in \mathbb{T}_{\tilde{u}, \tilde{i}}^+} \sum_{t^- \in \mathbb{T}_{\tilde{u}, \tilde{i}}^-} f(\hat{y}_{\tilde{u}, \tilde{i}, t^-} - \hat{y}_{\tilde{u}, \tilde{i}, t^+}) \right) \\ & + \alpha (\text{tr}(U^T L_U U) + \text{tr}(I^T L_I I) + \text{tr}(T^T L_T T)) \\ & + \beta (\|\mathcal{U}\|^2 + \|\mathcal{I}\|^2 + \|\mathcal{T}\|^2) \end{aligned} \quad (14)$$

⁶The user affinity graph and image affinity graph are constructed based on statistics of cojoined groups and visual similarity, respectively.

where $\|\mathcal{U}\|^2 + \|\mathcal{I}\|^2 + \|\mathcal{T}\|^2$ is l_1 regularization term to penalize large parameters, α and β are weights controlling the strength of

TABLE II
DOMINATING USER-SPECIFIC TOPICS FOR TWO EXAMPLE USERS

User A	Topic 1	military, aircraft , battleship, navy, artillery, iraq, aircraftcarrier, barracks
	Topic 2	apple, computer, art, girl, cellphone, cool, vintage, digital
	Topic 3	athlete, basketball, baseball, actor, sports, art, film, black

User B	Topic 1	buddha, buddhist, temple, religion, buddhism, thailand, asia, ancient
	Topic 2	blossoms, blooms, nature, macro, flower, bravo, butterfly, spring
	Topic 3	airplane, boeing, aircraft , airport, aviation, jet, aeroplane, cockpit

corresponding constraints. Obviously, directly optimizing (14) is infeasible and we employ an iterative optimization strategy.

IV. USER-SPECIFIC TOPIC MODELING

With the reconstructed user-tag-image ternary interrelations, we can directly perform the personalized image search: when user u submits a query q , the rank of image i is inversely proportional to the probability of u annotating i with tag q

$$\text{rank}(i|q, u) \propto \frac{1}{\hat{y}_{u,i,q}}. \quad (15)$$

However in practice, the queries and tags do not follow one-to-one relationship - one query usually corresponds to several related tags in the tag vocabulary. Besides, the query-tag correspondence differs from user to user. Therefore, we build topic spaces for each user to exploit this user-specific one-to-many relationship.

We investigate on a Flickr dataset of 270-K images that the average number of annotated images per user is only 30. The detailed distribution is shown in Fig. 4(b). Obviously the original annotation is too sparse to perform topic modeling, hence we employ the predicted annotations.⁷ Particularly, for each user u , the tags with 100 highest $\hat{y}_{u,i,t}$ are reserved as the annotations for image i . Each user's annotations to all the images constitute one *corpus*, and we choose Latent Dirichlet Allocation (LDA, [32]) to perform topic modeling. The individual tag is viewed as *word*, while the user's annotation to one image corresponds to one *document*.

LDA assumes that in one corpus, documents are generated from a set of K latent topics $\{\text{topic}_1, \dots, \text{topic}_K\}$. Document \mathbf{t}_i is the tags assigned to image i by individual user. In \mathbf{t}_i , each word t is associated with a latent topic. The generative process for user u 's annotation corpus \mathbb{D}_u is:

For each document \mathbf{t}_i in a corpus \mathbb{D}_u ,

- sample a K -vector document-topic distribution $\theta = p(\mathbf{topic}|\mathbf{t}_i, u)$ from a Dirichlet distribution;
- for each word t , sample topic assignment j according to θ and draw a word from the j th topic-word distribution $p(t|\text{topic}_j, u)$.

The generative model is fitted using a Gibbs sampler.

After the user-specific topic modeling, for each user u , we obtain 1) user-specific topics $p(t|\text{topic}_{1,\dots,K}^u)$ and 2) topic-sensitive user preference $p(\mathbf{topic}|\mathbf{t}_i, u)$, $i = 1, \dots, |\mathbb{D}_u|$.⁸ In Table II,

⁷Another reason is that individual user's tagging vocabulary is usually limited.

⁸For simplicity, we denote it as $p(\mathbf{topic}|i, u)$ in the following.

we list the first three dominating topics for two example users. Each topic is characterized by its eight most probable tags. The rank of the user u -specific topics is decided by $p(\text{topic}_j|u)$, the probability that user u is interested in topic j . This can be calculated by aggregating user u 's preference over all the images

$$p(\text{topic}_j|u) = \frac{\sum_i^{|\mathcal{I}|} p(\text{topic}_j|i, u)}{\sum_j^K \sum_i^{|\mathcal{I}|} p(\text{topic}_j|i, u)}. \quad (16)$$

From the user-specific topics, we can see:

- user's interest profile, e.g., user A is likely to be a military fan who also likes digital product and sports, while user B is keen at religion and interests in gardening and aircraft;
- the same tag may have different topic posterior distributions for different users, e.g., for user A , "aircraft" occurs frequently in a military-related topic, while for user B , "aircraft" returns to its literal sense of air vehicle.

A. Online Personalized Search

In the online stage, when user u submits a query q ,⁹ we first perform user-specific query mapping—estimate the conditional probability that q belongs to user u -specific topics

$$p(\text{topic}_j|q, u) = \frac{p(\text{topic}_j|u) \cdot p(q|\text{topic}_j, u)}{p(q)} \propto p(\text{topic}_j|u) \cdot p(q|\text{topic}_j, u) \quad j = 1, \dots, K. \quad (17)$$

From Table II, since user A has a principle interest on topic 1 ($p(\text{topic}_1|u)$) and "aircraft" has a high probability in topic 1 ($p(q|\text{topic}_1, u)$), when user A searches "aircraft", the query will have a high proportion on user A 's topic 1. The query distribution is then utilized as weights to compute user u 's topic-sensitive preferences over the images under the query q . The rank of image i can be obtained as

$$\text{rank}(i|q, u) \propto \frac{1}{\sum_j^K p(\text{topic}_j|q, u)p(\text{topic}_j|i, u)}. \quad (18)$$

When user A searches "aircraft", the images likely to be annotated by military-related tags are ranked higher according to (18). While, when user B searches "aircraft", the images likely to be annotated by aircraft-related tags will be ranked higher. We can see that the query relevance and user preference are simultaneously integrated into this personalized formulation.

⁹We only consider the one word-based queries in this paper and handling complex queries of multiple words is our future work.

TABLE III
DATASET STATISTICS FOR SEVERAL TAGGING SYSTEMS

Dataset	User $ \mathcal{U} $	Tag $ \mathcal{T} $	Document $ \mathcal{D} $	Observed Tagging Triplet $ \mathcal{O} $	Sparseness
Bibsonomy [13]	116	412	361	10,148	99.94%
Last.fm [34]	2,917	1,853	2,045	219,702	99.998%
NUS-WIDE	50,120	5,018	247,849	4,541,326	99.999991%
NUS-WIDE15	3,372	5,018	124,099	1,223,254	99.99994%

V. EXPERIMENTS

A. Dataset

We perform the experiments on a large-scale web image dataset, NUS-WIDE [33]. It contains 269 648 images with 5018 unique tags collected from Flickr. We crawled the images' owner information and obtained owner user ID of 247 849 images.¹⁰ The collected images belong to 50 120 unique users. Fig. 4 shows the distributions of #tagger and #tagged images.¹¹

We investigate on the sparseness of several social tagging systems in Table III, where *sparseness* is defined as

$$\text{sparseness} = 1 - \frac{|\mathcal{O}|}{|\mathcal{U}| \times |\mathcal{T}| \times |\mathcal{D}|}. \quad (19)$$

The results presented are not meant to be exhaustive but illustrative of the fact that Flickr has a more severe sparseness problem. We select the users owning no less than 15 images and keep their images to perform the tensor factorization, which is referred as NUS-WIDE15.

B. Parameter Setting

NUS-WIDE15 is randomly split into two parts, 90% for training and testing (denoted as \mathcal{S}), and 10% for validation (denoted as \mathcal{V}). The result of annotation prediction directly affect the performance of personalized search. In our work, we select parameters according to the performance of annotation prediction.¹² There are three sets of parameters for the proposed RMTF+LDA model. The first three parameters are the rank of factor matrices, r_U, r_I, r_T . According to [30] and [31], we simply choose the ranks proportional to the original dimensions $|\mathcal{U}|, |\mathcal{I}|, |\mathcal{T}|$ and set $r_U = 50, r_I = 250, r_T = 5$. This guarantee that the same ratio of energies for different modes are preserved. The second set of parameters are the regularization weights α, β . They control how much the tensor decomposition incorporates the information of affinity intra-relations. By keeping $r_U = 50, r_I = 250, r_T = 5$, we conduct a simple training of RMTF to choose α and β on the validation set. For each user, one post is randomly removed for annotation prediction evaluation. Fig. 5 illustrates the impacts of α and β on the F1 score of annotation prediction for top-10 recommended tags. We can see that the performance remain relatively steady when α and β change within a certain range. We set $\alpha = 0.01$ and $\beta = 0.001$, which achieves the highest average F1 score. The most important parameter for user-specific topic modeling is the number of latent topics for each user. For now the number

¹⁰Due to link failures, the owner ID of some images is unavailable.

¹¹The statistic of Del.icio.us is computed from our collected 50 000 webpages returned by the 20 most popular tags at <http://www.delicious.com/?view=hotlist>.

¹²The evaluation for annotation prediction is detailed in the next subsection.

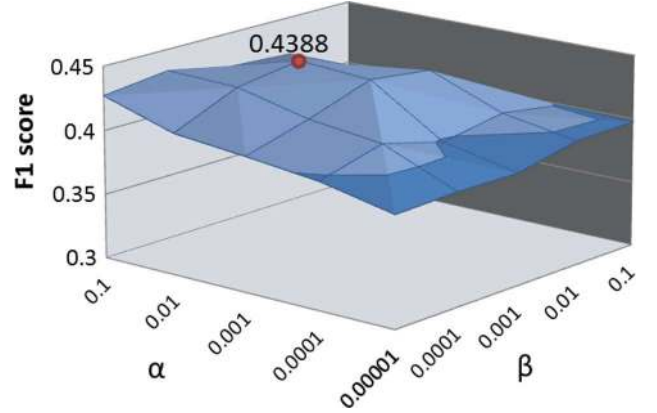


Fig. 5. Impact of parameters α and β .

is set same for different users and $K = 20$. We investigate the influence of K in the following experiment.

C. Annotation Prediction

We propose the novel RMTF model for users' annotation prediction. In this subsection, we first evaluate the performance of RMTF for annotation prediction. Following the evaluation process from [13], for each user we randomly remove all triplets he/she has annotated for one image to constitute the test set $\mathcal{S}_{\text{test}}$ —i.e., we remove one post for each user. The remaining observed user-image-tag triplets are used for regularized tensor factorization. Then we learn the model and predict top- N lists for each of the removed posts $\mathcal{P}_{\mathcal{S}_{\text{test}}}$ based on the reconstructed tensor $\hat{\mathcal{Y}}$ from (2). We compute the recall and precision of the top- N recommended tags and report the F1 score of the average recall and precision

$$\begin{aligned} \text{precision}(N) &:= \frac{1}{|\mathcal{S}_{\text{test}}|} \sum_{(u,i) \in \mathcal{P}_{\mathcal{S}_{\text{test}}}} \frac{|\text{Top}(u, i, N) \cap \{t | (u, i, t) \in \mathcal{S}_{\text{test}}\}|}{N} \end{aligned} \quad (20)$$

$$\begin{aligned} \text{recall}(N) &:= \frac{1}{|\mathcal{S}_{\text{test}}|} \sum_{(u,i) \in \mathcal{P}_{\mathcal{S}_{\text{test}}}} \frac{|\text{Top}(u, i, N) \cap \{t | (u, i, t) \in \mathcal{S}_{\text{test}}\}|}{|\{t | (u, i, t) \in \mathcal{S}_{\text{test}}\}|} \end{aligned} \quad (21)$$

$$\begin{aligned} \text{F1 score}(N) &:= \frac{2 \cdot \text{precision}(N) \cdot \text{recall}(N)}{\text{precision}(N) + \text{recall}(N)}. \end{aligned} \quad (22)$$

Four annotation prediction methods are performed for comparisons: most popular tags for image (Popular_I), most popular tags for user (Popular_U), FolkRank [13] and HOSVD [14].

TABLE IV
TESTING SET STATISTICS FOR EVALUATION

Testing set	# User	# Query	# Images tagged/favorited	# Tags annotated
NUS-WIDE15_A10_30	30	11	253	14,148
NUS-WIDE15_A100	30	18	4,566	319,702
NUS-WIDE15_F10_30	30	15	233	5,015
NUS-WIDE15_F100	19	15	3,214	19,254

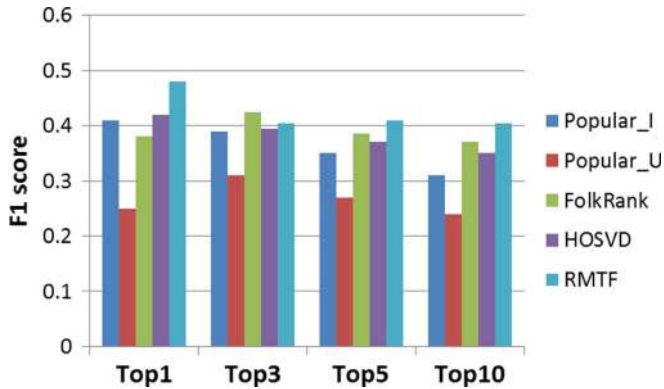


Fig. 6. F1 score of annotation prediction for different methods.

Fig. 6 illustrates the results. It is shown that RMTF generally performs the best, and with the increasing number of recommended tags, the F1 score decreases less steeper for RMTF than the other methods. This coincides with our discussions in the introduction that the proposed ranking scheme as well as exploiting the tag semantic-and-context relevance better alleviates the severe sparsity and noisy problem for Flickr dataset.

D. Personalized Search

In the research community of personalized search, evaluation is not an easy task since relevance judgement can only be evaluated by the searchers themselves. The most widely accepted approach is user study [24], [26], [17], [25], where participants are asked to judge the search results. Obviously this approach is very costly. In addition, a common problem for user study is that the results are likely to be biased as the participants know that they are being tested. Another extensively used approach is by user query logs or click-through history [20], [35]. However, this needs a large-scale real search logs, which is not available for most of the researchers.

Social sharing websites provide rich resources that can be exploited for personalized search evaluation. User's social activities, such as rating, tagging and commenting, indicate the user's interest and preference in a specific document. Recently, two types of such user feedback are utilized for personalized search evaluation. The first approach is to use social annotations [2], [11], [13]. The main assumption behind is that the documents tagged by user u with tag t will be considered relevant for the personalized query (u, t) . Another evaluation approach is proposed for personalized image search on Flickr [5], where the images marked *Favorite* by the user u are treated as relevant when u issues queries. The two evaluation approaches have their pros and cons and supplement for each other. We use both in our experiments and list the results in the following.

We select two state-of-the-art models as the baseline.

- Topic-based: topic-based personalized search using folksonomy [2].
- Preference-based: personalized image search by predicting user interests-based preference [5].

Note that both methods follow the two-step scheme: the overall ranking is decided by separately computing query relevance and user preference. In addition, we also compared the performances of the proposed model with different settings.

- TF_0/1_LDA: TF without smoothness constraints, optimization under the *0/1 scheme*, using user-specific topic modeling.
- MTF_0/1_LDA: TF with multicorrelation smoothness constraints, optimization under the *0/1 scheme*, using user-specific topic modeling.
- RMTF_LDA, the proposed model: annotations predictions by RMTF, using user-specific topic modeling.
- RMTF: Directly using the RMTF-based predicted annotations for personalized rank according to (15).

1) *Annotation-Based Evaluation*: We follow Xu's [2] evaluation framework and first compare the performances according to users' original annotations. To perform the evaluation in the situations of users with different amount of original annotations, we build two test scenarios: 1) 30 randomly selected users who tagged 10–30 images and their tagging records, denoted as NUS-WIDE15_A10_30. 2) 30 randomly selected users who tagged more than 100 images and their tagging records, denoted as NUS-WIDE15_A100. For NUS-WIDE15_A100, the overlapping 18 tags the 50 users used are selected as the test queries, while for NUS-WIDE15_A10_30, the number of test queries is 11. The statistics of the testing sets are shown in Table IV. In order to reduce the dependency between original annotations and evaluation, we remove the tagging data related to the test queries. It is done as follows: for each personal query (u, t) , we remove the triplets $(u, :, t)$ from the training set.

In the experiments, we use average precision (AP) as the evaluation metric, which is a widely used relevance metric evaluating the performance of the top documents in the ranked list. AP is defined as

$$AP = \frac{\sum_{i=1}^{n^+} \sum_{j=1}^i \frac{r(j)}{i} * r(i)}{\sum_{i=1}^{n^+} r(i)} \quad (23)$$

where i is the position (i.e., rank) of the document, n^+ is the number of the relevant documents and $r(i)$ denotes the relevance of the document in position i . In our case, a binary value for $r(i)$ is used by setting it to 1 if the document is relevant and 0 otherwise. The Mean Average Precision (MAP) is the mean of




Fig. 7. mMAP value of personalized search for different methods (a) Annotation-based. (b) Favorite-based.

the APs of all queries for one user, and the mean of the MAPs of all test users is referred as mMAP

$$\text{mMAP} = \frac{\sum_{i=1}^{N_u} \frac{\sum_{j=1}^{N_q} AP(i,j)}{N_q}}{N_u}. \quad (24)$$

Here $AP(i, j)$ represents the AP value of the j th query for the i th user, and N_u, N_q is the number of test users and queries, respectively.

The results are shown in Fig. 7(a). *Non-personalized* denotes the non-personalized rank result by only considering the query relevance. We can see that all the personalized methods outperform the non-personalized scheme. Comparing between the two test scenarios of NUS-WIDE15_A10_30 and NUS-WIDE15_A100, the performances of personalized methods improve as the test users' original annotations increase. This is reasonable as these methods utilize the social annotation resources and the more user feedback is available, the more accurate user preferences can be estimated. What is interesting is that the preference-based model [5] and the proposed model are more sensitive to the amount of original annotations. The reason may be that [5] and our methods extract topic spaces by explicitly exploiting the tagging data, while in the topic-based model [2], the topic space is predefined and the original annotation is just used to generate the topic vector.

Focusing on either test scenario, the performance of the proposed RMTF_LDA, even MTF_0/1_LDA, is superior than the baseline methods, which demonstrate the advantage of simultaneously considering query relevance and user preference over the separate schemes. Depending on one-to-one query-tag assumption, the performance of RMTF deteriorates dramatically without the user-specific topic modeling. Moreover, RMTF_LDA outperforms MTF_0/1_LDA, showing the advantage of the proposed ranking scheme over the conventional 0/1 scheme. Without smoothness priors, TF_0/1 fails to preserve the affinity structures and achieves inferior results.

2) *Favorite-Based Evaluation*: There is a delicate issue with annotation-based evaluation. Both the input to the personalized models and the evaluation for the output results are based on the original annotations. Although the specific tagging data $(u, :, t)$ have been removed when testing the personal query (u, t) , as individual user's tagging vocabulary tends to be limited, the

remaining annotations will implicitly provide the association between u and t . For example, assuming one user u usually tag "wildlife" and "animal" together, when he/she issues "wildlife" as test query, though all $(u, :, \text{wildlife})$ have been removed from the training process, regarding "wildlife" and "animal" are likely to have a close relation in the user-specific topics, the images tagged by "animal" will be given high probability and guide the final rank. On Flickr, users are encouraged to express their preference on images by adding *Favorite* marks. Illustrated by Lu's evaluation framework [5], we employ users' *Favorite* marks for evaluation, which are not used in the training process. This guarantees that personalization is evaluated without any prior knowledge.

To be consistent with the annotation-based evaluation scheme, we also build two test scenarios for the favorite-based evaluation: 1) 30 randomly selected users who added *Favorite* to 10–30 images, denoted as NUS-WIDE15_F10_30, and 2) all the 19 users who added *Favorite* to more 100 images, denoted as NUS-WIDE15_F100. 15 tags frequently appearing in the annotation of those favorite images are selected as the test queries. The metric of mMAP is utilized to evaluate the performance and the results are demonstrated in Fig. 7(b). We have the following observations.

- The mMAP is relatively low compared with Annotation-based evaluation. This phenomenon reflects the problem of Favorite-based evaluation scheme: the *Favorite* images are considered relevant for all the test queries. As no query information is involved, for those queries nonrelevant with the topic of the *Favorite* images, the AP tends to be low;
- Comparing between the two test scenarios, the average performance of NUS-WIDE15_F100 also improves over NUS-WIDE15_F10_30, but not as significant as in Annotation-based evaluation. One possible reason for the improvement is that those users having more *Favorite* marks are active users who are likely to also attend more interest groups and tag more images. While, the improvement is not so significant demonstrates that the Favorite-based evaluation scheme is less sensitive to the amount of original annotations;
- Another obvious difference from the results of annotation-based evaluation is that the performance of TF_0/1 and MTF_0/1_LDA degrade dramatically. The mMAP of TF_0/1 is even lower than the non-personalized method.

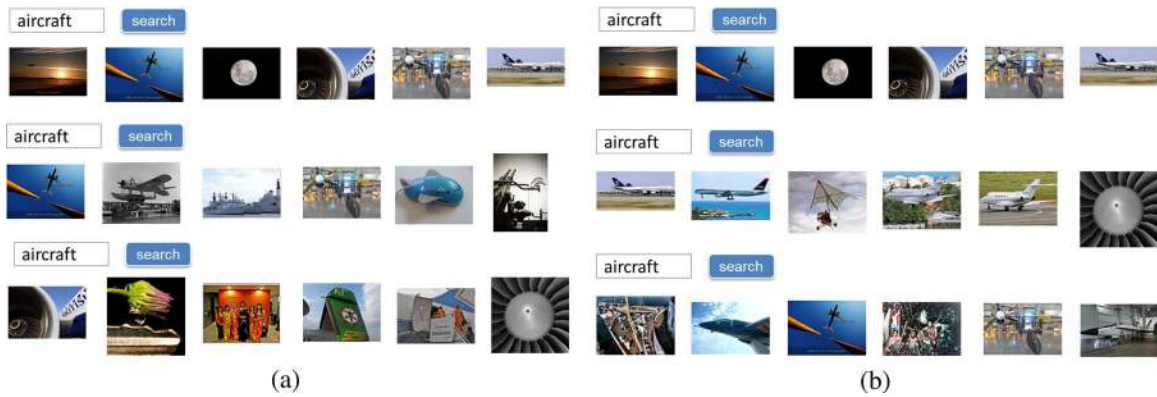


Fig. 8. Example of non-personalized (*top*) and personalized (*middle* for User A and *bottom* for User B) search results for query “aircraft” (a) Topic-based method; (b) RMTF_LDA.

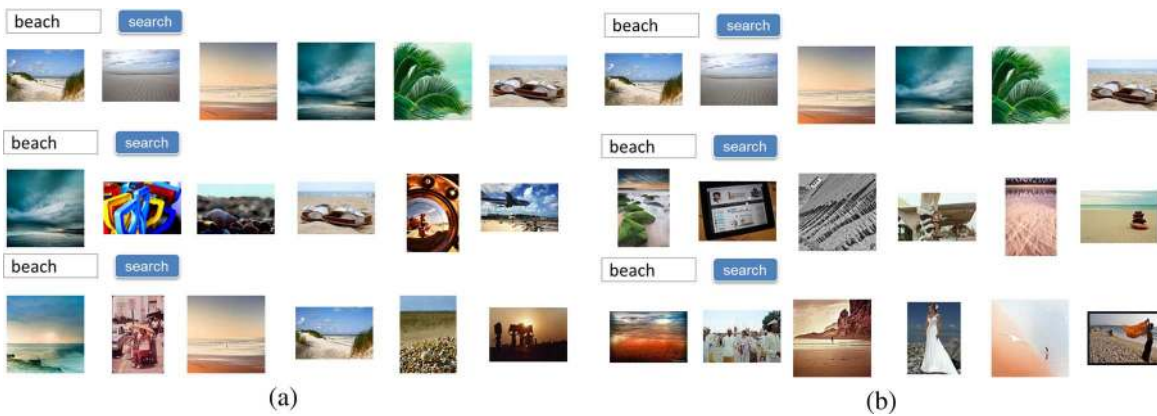


Fig. 9. Example of non-personalized (*top*) and personalized (*middle* for User A and *bottom* for User B) search results for query “beach.” (a) Topic-based method. (b) RMTF_LDA.

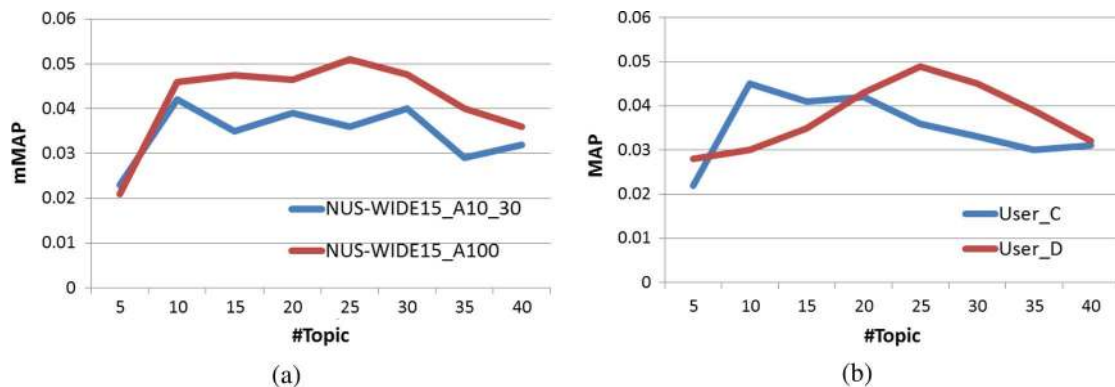


Fig. 10. Influence of topic number K (a) mMAP for all test users; (b) MAP for two users from NUS-WIDE15-A100.

For the annotation-based evaluation, TF_{0/1} achieves comparable results due to the implicit prior knowledge provided by the original annotations. By utilizing the *Favorite* marks, a heterogeneous resource for evaluation, the implicit prior is eliminated.

Fig. 8 displays exemplary search results for the query “aircraft”. The top six non-personalized results and the personalized results of User A and User B mentioned in Section IV-A are shown. We can see that by simultaneously considering query relevance and user information, the proposed RMTF_LDA captures the user’s preference under certain topics. As a result of mapping “aircraft” to Topic 1 of Table II, the top search re-

sults for user A mainly focus on aerocrafts. While, for user B, the top search results are basically military related, which coincides with user B’s preference. For the baseline method which separate query relevance and user preference, sometimes the search results are hard to interpret. For example, the second and third images for user B in Fig. 8(a) are ranked higher because user B has a major interest in religion and flower. However, these images have little relation with aircraft. We note that for some general queries which have clear search intents, personalized search tends to fail. Fig. 9 illustrates one of such examples. With “beach” having common understanding to different users, incorporating user information will generate con-

fusing search results. There are literatures [35], [36] discussing the issue about when to perform personalization. It seems that the benefit of personalization is highly dependent on the ambiguity of the query. Since there is no conclusion to this problem, in this paper we focus on the problem of how to perform personalization and discussion of when to perform personalization is beyond the scope of this paper.

3) *Influence of Topic Number K* : For the standard LDA, the number of latent topics needs to be specified. In the above experiments, we set the same number of topics for all users and $K = 20$. In the following, we vary the selection of K and investigate the influence of topic number. We utilize the annotation-based evaluation scheme to compute mMAP. The results are illustrated in Fig. 10. It is shown that mMAP displays no definite trend as K changes [Fig. 10(a)], while for individual user C and user D there exist obvious optimal K (see Fig. 10(b); obviously user C and D have an optimal $K = 10$ and $K = 25$, respectively). This observation is inline with the expectations that users have different topic spaces and validate the necessity of user-specific topic modeling. In addition, user-specific topic number should be specified in the future work. There is a number of extension work on standard LDA to automatically select the number of topics. The most common one is HDP-LDA, which uses hierarchical Dirichlet processes (HDP, [37]) to model the Dirichlet mixtures in LDA nonparametrically.

VI. CONCLUSION AND FUTURE WORK


How to effectively utilize the rich user metadata in the social sharing websites for personalized search is challenging as well as significant. In this paper we propose a novel framework to exploit the users' social activities for personalized image search, such as annotations and the participation of interest groups. The query relevance and user preference are simultaneously integrated into the final rank list. Experiments on a large-scale Flickr dataset show that the proposed framework greatly outperforms the baseline.

In the future, we will improve our current work along four directions. 1) In this paper, we only consider the simple case of one word-based queries. Actually, the construction of topic space provides a possible solution to handle the complex multiple words-based queries. We will leave it for our future work. 2) During the user-specific topic modeling process, the obtained user-specific topics represent the user's distribution on the topic space and can be considered as user's interest profile. Therefore, this framework can be extended to any applications based on interest profiles. 3) For batch of new data (new users or new images), we directly restart the RMFTF and user-specific topic modeling process. While, for a small amount of new data, designing the appropriate update rule is another future direction. 4) Utilizing large tensors brings challenges to the computation cost. We plan to turn to parallelization (e.g., parallel MATLAB) to speedup the RMFTF converge process. Moreover, the distributed storing mechanism of parallelization will provide a convenient way to store very large matrices and further reduce the storage cost.

REFERENCES

- [1] B. Smyth, "A community-based approach to personalizing web search," *Computer*, vol. 40, no. 8, pp. 42–50, 2007.
- [2] S. Xu, S. Bao, B. Fei, Z. Su, and Y. Yu, "Exploring folksonomy for personalized search," in *Proc. SIGIR*, 2008, pp. 155–162.
- [3] D. Carmel, N. Zwerdling, I. Guy, S. Ofek-Koifman, N. Har'El, I. Ronen, E. Uziel, S. Yogev, and S. Chernov, "Personalized social search based on the user's social network," in *Proc. CIKM*, 2009, pp. 1227–1236.
- [4] Y. Cai and Q. Li, "Personalized search by tag-based user profile and resource profile in collaborative tagging systems," in *Proc. CIKM*, 2010, pp. 969–978.
- [5] D. Lu and Q. Li, "Personalized search on flickr based on searcher's preference prediction," in *Proc. WWW*, 2011, pp. 81–82, companion volume.
- [6] P. Heymann, G. Koutrika, and H. Garcia-Molina, "Can social bookmarking improve web search?," in *Proc. WSDM*, 2008, pp. 195–206.
- [7] S. Bao, G.-R. Xue, X. Wu, Y. Yu, B. Fei, and Z. Su, "Optimizing web search using social annotations," in *Proc. WWW*, 2007, pp. 501–510.
- [8] D. Zhou, J. Bian, S. Zheng, H. Zha, and C. L. Giles, "Exploring social annotations for information retrieval," in *Proc. WWW*, 2008, pp. 715–724.
- [9] J. Tang, S. Yan, R. Hong, G. Qi, and T. Chua, "Inferring semantic concepts from community-contributed images and noisy tags," in *Proc. ACM Multimedia*, 2009, pp. 223–232.
- [10] J. Tang, H. Li, G. Qi, and T. Chua, "Image annotation by graph-based inference with integrated multiple/single instance representations," *IEEE Trans. Multimedia*, vol. 12, no. 2, pp. 131–141, Feb. 2010.
- [11] M. J. Carman, M. Baillie, and F. Crestani, "Tag data and personalized information retrieval," in *Proc. SSM*, 2008, pp. 27–34.
- [12] R. Jäschke, L. B. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme, "Tag recommendations in folksonomies," in *Proc. PKDD*, 2007, pp. 506–514.
- [13] R. Jäschke, L. B. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme, "Tag recommendations in social bookmarking systems," *AI Commun.*, vol. 21, no. 4, pp. 231–247, 2008.
- [14] P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos, "A unified framework for providing recommendations in social tagging systems based on ternary semantic analysis," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 2, pp. 179–192, Feb. 2010.
- [15] G. Zhu, S. Yan, and Y. Ma, "Image tag refinement towards low-rank, content-tag prior and error sparsity," in *Proc. ACM Multimedia*, 2010, pp. 461–470.
- [16] J. Teevan, M. R. Morris, and S. Bush, "Discovering and using groups to improve personalized search," in *Proc. WSDM*, 2009, pp. 15–24.
- [17] P.-A. Chirita, W. Nejdl, R. Paiu, and C. Kohlschütter, "Using odp metadata to personalize search," in *Proc. SIGIR*, 2005, pp. 178–185.
- [18] R. Kraft, F. Maghoul, and C.-C. Chang, "Y!q: Contextual search at the point of inspiration," in *Proc. CIKM*, 2005, pp. 816–823.
- [19] K. Sugiyama, K. Hatano, and M. Yoshikawa, "Adaptive web search based on user profile constructed without any effort from users," in *Proc. WWW*, 2004, pp. 675–684.
- [20] J.-T. Sun, H.-J. Zeng, H. Liu, Y. Lu, and Z. Chen, "Cubesvd: A novel approach to personalized web search," in *Proc. WWW*, 2005, pp. 382–390.
- [21] F. Qiu and J. Cho, "Automatic identification of user interest for personalized search," in *Proc. WWW*, 2006, pp. 727–736.
- [22] Y. Liu, T. Mei, and X. Hua, "CrowdReranking: Exploring multiple search engines for visual search reranking," in *Proc. SIGIR*, 2009, pp. 500–507.
- [23] N. D. Lane, D. Lymberopoulos, F. Zhao, and A. T. Campbell, "Hapori: Context-based local search for mobile phones using community behavioral modeling and similarity," in *Proc. UbiComp*, 2010, pp. 109–118.
- [24] J. E. Pitkow, H. Schütze, T. A. Cass, R. Cooley, D. Turnbull, A. Edmonds, E. Adar, and T. M. Breuel, "Personalized search," *Commun. ACM*, vol. 45, no. 9, pp. 50–55, 2002.
- [25] P.-A. Chirita, C. S. Firan, and W. Nejdl, "Personalized query expansion for the web," in *Proc. SIGIR*, 2007, pp. 7–14.
- [26] J. Teevan, S. T. Dumais, and E. Horvitz, "Personalizing search via automated analysis of interests and activities," in *SIGIR*, 2005, pp. 449–456.
- [27] K. Lerman, A. Plangprasopchok, and C. Wong, "Personalizing image search results on flickr," *J. CoRR*, vol. abs/0704.1676, 2007.
- [28] W.-J. Li and D.-Y. Yeung, "Relation regularized matrix factorization," in *Proc. IJCAI*, 2009, pp. 1126–1131.

- [29] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999, pp. 02178–9998.
- [30] H. A. L. Kiers and A. D. Kinderen, “A fast method for choosing the number of components in Tucker3 analysis,” *British J. Math. Statistical Psychol.*, vol. 56, pp. 119–125, 2003.
- [31] T. G. Kolda and B. W. Bader, “Tensor decompositions and applications,” *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, 2009.
- [32] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent Dirichlet allocation,” *J. Mach. Learning Res.*, vol. 3, pp. 993–1022, 2003.
- [33] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y. Zheng, “Nus-wide: A real-world web image database from National University of Singapore,” in *Proc. CIVR*, 2009, pp. 1–9.
- [34] S. Rendle, “Factorization machines,” in *Proc. ICDM*, 2010, pp. 995–1000.
- [35] Z. Dou, R. Song, and J.-R. Wen, “A large-scale evaluation and analysis of personalized search strategies,” in *Proc. WWW*, 2007, pp. 581–590.
- [36] J. Teevan, S. T. Dumais, and D. J. Liebling, “To personalize or not to personalize: Modeling queries with variation in user intent,” in *Proc. SIGIR*, 2008, pp. 163–170.
- [37] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, “Sharing clusters among related groups: Hierarchical Dirichlet processes,” in *Proc. NIPS*, 2004, pp. 1–8.



Jitao Sang received the B.E. degree from the South-East University, Nanjing, China, in 2007, and is currently pursuing the Ph.D. degree at the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China.

In 2010 and 2011, he was an intern student in the China-Singapore Institute of Digital Media (CSIDM) and Microsoft Research Asia (MSRA), respectively. His current research interests include multimedia content analysis, social media mining, computer vision and pattern recognition.



Changsheng Xu (M'97-SM'99) received the Ph.D from Tsinghua University, Beijing, China, in 1996.

He is a Professor at the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China, and Executive Director of the China-Singapore Institute of Digital Media, Singapore. He has held 30 granted/pending patents and published over 200 refereed research papers. His research interests include multimedia content analysis/indexing/retrieval, pattern recognition and computer vision.

Dr. Xu received the Best Editorial Member Award in 2008. He is an Associate Editor of *ACM Transactions on Multimedia Computing, Communication and Applications* and served as Program Chair of *ACM Multimedia 2009*. He has served as Associate Editor, Guest Editor, General Chair, Program Chair, Area/Track Chair, Special Session Organizer, Session Chair, and TPC member for over 20 IEEE and ACM prestigious multimedia journals, conferences and workshops.



Dongyuan Lu received the B.S. degree from Beijing Normal University, Beijing, China, in 2007, and is currently pursuing the Ph.D. degree at the State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing, China.

Her research interests include social network analysis, web data mining, and information retrieval.