

## LA-CTR: A Limited Attention Collaborative Topic Regression for Social Media

**Jeon-Hyung Kang**

USC Information Sciences Institute  
4676 Admiralty Way  
Marina del Rey, CA 90292  
jeonhyuk@usc.edu

**Kristina Lerman**

USC Information Sciences Institute  
4676 Admiralty Way  
Marina del Rey, CA 90292  
lerman@isi.edu

### Abstract

Probabilistic models can learn users' preferences from the history of their item adoptions on a social media site, and in turn, recommend new items to users based on learned preferences. However, current models ignore psychological factors that play an important role in shaping online social behavior. One such factor is attention, the mechanism that integrates perceptual and cognitive features to select the items the user will consciously process and may eventually adopt. Recent research has shown that people have finite attention, which constrains their online interactions, and that they divide their limited attention non-uniformly over other people. We propose a collaborative topic regression model that incorporates limited, non-uniformly divided attention. We show that the proposed model is able to learn more accurate user preferences than state-of-art models, which do not take human cognitive factors into account. Specifically we analyze voting on news items on the social news aggregator and show that our model is better able to predict held out votes than alternate models. Our study demonstrates that psycho-socially motivated models are better able to describe and predict observed behavior than models which only consider latent social structure and content.

### Introduction

The rise of social media has drastically exacerbated the problem of information overload. Social media users create far more information than other users can process. Social media sites address this problem by allowing users to follow updates from specified friends; however, as the number of friends a user follows grows over time, information overload returns. One solution to this problem is to recommend items to a user based on his or her interests or preferences, which can be learned by observing user's activity. Researchers have developed statistical models to learn users' preferences from the items they, and their friends, have adopted in social media (Chua, Lauw, and Lim 2012; Purushotham, Liu, and Kuo 2012), and filter the streams of new information accordingly.

Existing statistical models largely ignore the psychological and social factors that shape user behavior. However, recent research has demonstrated that such factors, in particular attention, play an important role in social media users'

behavior and interactions (Counts and Fisher 2011; Hodas and Lerman 2012; Goncalves, Perra, and Vespignani 2011; Weng et al. 2012). Attention is the mechanism that integrates perceptual and cognitive factors to select the small fraction of input to be processed in real time (Kahneman 1973; Rensink, O'Regan, and Clark 1997). Attentive acts, such as reading a tweet, browsing a web page, or responding to email, require mental effort, and since the brain's capacity for mental effort is limited, so is attention. Limited attention was shown to constrain user's social interactions (Goncalves, Perra, and Vespignani 2011) and the spread of information (Hodas and Lerman 2012) in social media. Moreover, while users divide their limited attention over their friends (Hodas and Lerman 2012), some friends receive a larger share of their attention than others (Gilbert and Karahalios 2009).

In this paper, we close the gap between behavioral and statistical recommendation models by presenting  $\mathcal{L}A$ -CTR, a model that extends collaborative recommendation topic model CTR introduced by Wang & Blei (Wang and Blei 2011) by integrating users' limited attention into it. Like CTR,  $\mathcal{L}A$ -CTR uses item adoptions of all users and the content of items as a basis for its recommendations. However,  $\mathcal{L}A$ -CTR modifies the existing model in two important ways. First, the model captures a user's limited attention, which he or she must divide among all others she follows. Second, the model includes a notion of influence, that is the fact that the user may preferentially pay more attention to some of the users, e.g., close friends (Gilbert and Karahalios 2009). Moreover, the model allows the influence to vary according to each user's topical interests. We evaluate the model on real world data from the social news site Digg. We demonstrate that the proposed psycho-socially motivated model performs better on the news story recommendation task than alternative models.

In the rest of the paper we first review existing research, which includes two state-of-art probabilistic collaborative models: CTR and CTR-smf. In Section " $\mathcal{L}A$ -CTR", we introduce a limited attention collaborative regression model that takes into consideration the limited, non-uniformly divided attention of social media users. In Section "Evaluation", we study  $\mathcal{L}A$ -CTR model with votes on news items on the social news aggregator Digg and show that our proposed model is better able to predict held out votes than alterna-

tive models that do not take limited attention into account. Our study demonstrates that behaviorally motivated models are better able to describe observed user behavior in social media, and lead to better recommendation tools.

### Related Work

The growing abundance of social media data has allowed researchers to start asking questions about the nature of influence, how social ties interact with the contents of items, and how these affect transmission of items in social networks (Bakshy et al. 2011; Romero, Meeder, and Kleinberg 2011). Recently, researchers demonstrated the importance of attention in online interactions (Goncalves, Perra, and Vespignani 2011; Hodas and Lerman 2012; Weng et al. 2012). Specifically relevant to this paper are studies that show that limited attention constrains social interactions (Goncalves, Perra, and Vespignani 2011) and information diffusion (Hodas and Lerman 2012) in social media. The model proposed in this paper incorporates these ideas into a hidden topic model for collaborative recommendation.

In text mining applications, hidden topic models, such as Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003), examine documents to identify hidden topics, each represented as distributions over groups of related words. Traditional topic models have been extended to a networked setting to model hyperlinks between documents (Nallapati and Cohen 2008), topical relations between them (Chang 2009; Daumé 2009), and the varying vocabularies and styles of different authors (Rosen-Zvi et al. 2004). More generally, researchers have applied topic models to social media analysis, e.g., by treating users as documents and content they create as words (Ramage et al. 2009). Then, by analyzing messages posted by many users, one can separate users into topical groups.

Collaborative filtering methods examine item recommendations of many people to discover their preferences and recommend new items that were liked by similar people. User-based (Herlocker et al. 1999) and item-based (Sarwar et al. 2001; Karypis 2000) recommendation approaches have been developed for predicting how people will rate new items. Matrix factorization-based latent-factor models (Salakhutdinov and Mnih 2008; Koren, Bell, and Volinsky 2009) have shown promise in creating better recommendations by incorporating user interests into the model. However all these approaches ignore the content of items in the recommendation task. Recently, probabilistic topic modeling techniques were merged with collaborative filtering to improve the explanatory power of recommendation tools (Agarwal and Chen 2010; Wang and Blei 2011). Content analysis based on probabilistic topic modeling has been proposed to incorporate into collaborative filtering (Agarwal and Chen 2010) to account for user rating and item popularity biases in the data. Authors show better prediction performance by regularization of both user and item factors through user features and the content of item. Collaborative Topic Regression (CTR) model (Wang and Blei 2011) incorporates collaborative filtering model based on probabilistic topic modeling. Both models do a good job in using item content for recommendation; however, neither takes social structure of users into

account.

Researchers have extended LDA to a social recommendation setting (Chua, Lauw, and Lim 2012), in which a user is likely to adopt an item based on the latent interests of other users. Social correlations between users are learned from their item adoption choices, rather than specified explicitly through the follower graph. Social recommendation system has been proposed by matrix factorization techniques for both user’s social network and their item rating histories (Ma et al. 2008; Yang et al. 2011) without using the contents of items. In CTR-smf (Purushotham, Liu, and Kuo 2012), authors integrate CTR with social matrix factorization models to take social correlation between users into account. However, these works utilize *homophily* effect in social media to smooth the similarity of users’ interests to their friends, instead of directly learning how users allocate their attention over their friends. So depending on the degree of similarity among connected users, social matrix factorization may not consistently perform well. The model proposed in this paper directly learns how much attention users allocate to other users, and leverages these learned influences to make better recommendations.

Attention in online interactions has been modeled in a previous work (Kang, Lerman, and Getoor 2013) to directly learn users’ limited attention over their friends and topics. However, that model was applied to dyadic data (Agarwal and Chen 2009; Koren, Bell, and Volinsky 2009; Yu et al. 2009) without providing any explanation about the learned topics. Furthermore, authors assumed a simple user item adoption mechanism, where item adoption is decided based on sampled interest and topic pairs, rather than the inner product of user interest and item topic latent vectors, as matrix factorization does. The model proposed in this paper makes two important advances. First, it incorporates the content of items, therefore, being able to provide explanatory power to its recommendations, as well as recommend previously unseen items. Furthermore, the proposed model is able to learn the degree of influence of other users in modeling limited attention.




Figure 1: The  $\mathcal{L}_A$ -CTR model.

## LA-CTR: Limited Attention Collaborative Topic Regression Model

We introduce a limited attention collaborative topic regression model (LA-CTR) that learns user preferences from the observations of items they adopt on a social media site. Our model captures the behavior of social media users by introducing three novel elements that model users' limited attention ( $\phi$ ), their interests ( $u$ ), and how much attention they pay to other users ( $s$ ), which we also refer to as influence. Two additional elements in the model are the item's topic for explaining recommendations ( $v$ ) and for explaining the item's contents ( $\theta$ ), which were introduced by (Wang and Blei 2011). Finally ratings ( $r$ ) are observed variables. Rating is equal to one when user adopts an item and zero otherwise.

Figure 1 presents LA-CTR in graphical form. Each variable is drawn from a normal distribution defined as follows:

$$\begin{aligned} u_i &\sim \mathcal{N}(0, \lambda_u^{-1} I_K) \\ s_i &\sim \mathcal{N}(0, \lambda_s^{-1} I_N) \\ \phi_{il} &\sim \mathcal{N}\left(g_\phi(s_{il}u_i), c_{il}^\phi \lambda_\phi^{-1} I_K\right) \\ r_{ijl} &\sim \mathcal{N}\left(g_r(\phi_{il}^T v_j), c_{ijl}^r\right) \end{aligned} \quad (1)$$

where  $N$  is the number of users,  $D$  is the number of items, and  $K$  is the number of topics. Symbol  $T$  refers to the transpose operation. We define  $g_\phi$  and  $g_r$  as linear functions for simplicity. The precision parameters  $c_{il}^\phi$  and  $c_{ijl}^r$  serve as confidence parameters for influence  $s_{il}$  and rating  $r_{ijl}$ . The confidence parameter  $c_{il}^\phi$  represents how much attention user  $i$  pays to user  $l$ . In social media, users adopt new items primarily based on their friends' recommendations, i.e., by seeing the items their friends adopt. However, users may also be exposed to items from outside of their social network, depending on the site's interface. We use  $c_{il}^\phi$  to model the structure of the social network: we set it a high value  $a^\phi$  when user  $l$  is a friend of user  $i$  and a low value  $b^\phi$  when he is not ( $a^\phi > b^\phi > 0$ ). Similarly, when user  $i$  does not adopt item  $j$  (i.e., her rating of item  $j$  is zero), it can be interpreted into two ways: either user  $i$  was aware of item  $j$  but did not like it, or user  $i$  would have liked item  $j$  but never saw it (Wang and Blei 2011). We set  $c_{ijl}^r$  as a confidence parameter for user  $i$ 's rating on item  $j$  via user  $l$  and set it to a high value  $a^r$  when  $r_{ijl} = 1$  and a low value  $b^r$  when  $r_{ijl} = 0$  ( $a^r > b^r > 0$ ). In this paper, we use the same confidence parameter values,  $a^r = a^\phi = 1.0$  and  $b^r = b^\phi = 0.01$ , for all  $c_{il}^\phi$  and  $c_{ijl}^r$ .

Following CTR, we use the hidden topic model LDA (Blei, Ng, and Jordan 2003) to capture item's topic distribution  $\theta$ , which is represented as a  $K$  dimensional vector. In LDA, the topic distribution ( $\theta_{1:D}$ ) of each document is viewed as a mixture of multiple topics, with each topic ( $\beta_{1:K}$ ) a distribution over words from a vocabulary of size  $M$ . Like CTR, we assume a generative process in which documents are created by drawing words from their topic distributions. Latent variable  $\epsilon_j \sim \mathcal{N}(0, \lambda_v^{-1} I_K)$  captures the differences between topics that explain the contents of documents and those that explain recommendations. Depending

on the choice of  $\lambda_v$ , item  $j$ 's topic distribution  $\theta_j$  is perturbed to create latent vector  $v_j$ , which could be similar to  $\theta_j$  or diverge from it:

$$v_j \sim \mathcal{N}(\theta_j, \lambda_v^{-1} I_K) \quad (2)$$

The generative process for item adoption through a social network can be formalized as follows:

For each user  $i$   
 Generate  $u_i \sim \mathcal{N}(0, \lambda_u^{-1} I_K)$   
 Generate  $s_i \sim \mathcal{N}(0, \lambda_s^{-1} I_N)$   
 For each user  $l$   
 Generate  $\phi_{il} \sim \mathcal{N}(g_\phi(s_{il}u_i), c_{il}^\phi \lambda_\phi^{-1} I_K)$   
 For each item  $j$   
 Generate  $\theta_j \sim \text{Dirichlet}(\alpha)$   
 Generate  $\epsilon_j \sim \mathcal{N}(0, \lambda_v^{-1} I_K)$  and set  $v_j = \epsilon_j + \theta_j$   
 For each word  $w_{jm}$   
 Generate topic assignment  $z_{jm} \sim \text{Mult}(\theta)$   
 Generate word  $w_{jm} \sim \text{Mult}(\beta_{z_{jm}})$   
 For each user  $i$   
 For each attention friend  $l$   
 For each adopted item  $j$   
 Choose the rating  $r_{ijl} \sim \mathcal{N}(\phi_{il}^T v_j, c_{ijl}^r)^{-1}$

Here  $\lambda_u = \sigma_r^2 / \sigma_u^2$ ,  $\lambda_v = \sigma_r^2 / \sigma_v^2$ ,  $\lambda_s = \sigma_r^2 / \sigma_s^2$ , and  $\lambda_\phi = \sigma_r^2 / \sigma_\phi^2$ . Note that latent vectors  $u$ ,  $v$  and  $\phi$  are in a shared  $K$ -dimensional space.

### Learning Parameters

To learn model parameters, we follow the approaches of CTR (Wang and Blei 2011) and CTR-smf (Purushotham, Liu, and Kuo 2012) and develop an EM-style algorithm to calculate the maximum a posteriori (MAP) estimates. MAP estimation is equivalent to maximizing the complete log likelihood ( $\ell$ ) of  $U$ ,  $S$ ,  $\phi$ ,  $V$ ,  $R$ , and  $\theta$  given  $\lambda_u$ ,  $\lambda_s$ ,  $\lambda_\phi$ ,  $\lambda_v$  and  $\beta$ . We set Dirichlet prior  $\alpha$  to 1.

$$\begin{aligned} \ell = & -\frac{\lambda_u}{2} \sum_i u_i^T u_i - \frac{\lambda_v}{2} \sum_j (v_j - \theta_j)^T (v_j - \theta_j) \\ & + \sum_j \sum_t \log \left( \sum_k \theta_{jk} \beta_{k, w_{jt}} \right) - \frac{\lambda_s}{2} \sum_i s_i^T s_i \\ & - \sum_i \sum_j \sum_l \frac{c_{ijl}^r}{2} (r_{ijl} - \phi_{il}^T v_j)^2 \\ & - \frac{\lambda_\phi}{2} \sum_i \sum_l c_{il}^\phi (\phi_{il} - s_{il} u_i)^T (\phi_{il} - s_{il} u_i) \end{aligned} \quad (3)$$

We use gradient ascent to estimate MAP and iteratively optimize the variables  $\{u_i, v_j, \phi_{il}, s_i\}$  and the topic proportions  $\theta_j$ . Given a current estimate, we take the gradient of  $\ell$  with respect to  $u_i, v_j, s_i$ , and  $\phi_{il}$  and set it to zero. Derived update

equations are following:

$$\begin{aligned}
u_i &\leftarrow \left( \lambda_u I_k + \lambda_\phi S_i^T C_i^\phi S_i \right)^{-1} \lambda_\phi \phi_i C_i^\phi S_i \\
v_j &\leftarrow \left( \lambda_v I_k + \phi C_j^r \phi^T \right)^{-1} \left( \phi C_j^r R_j + \lambda_v \theta_j \right) \\
S_i &\leftarrow \left( \lambda_s I_N + \lambda_\phi u_i^T C_i^\phi u_i \right)^{-1} \lambda_\phi C_i^\phi \phi_i^T u_i \\
\phi_{il} &\leftarrow \left( \lambda_\phi C_{il}^\phi I_K + V C_{il}^r V^T \right)^{-1} \left( V C_{il}^r R_{il} + \lambda_\phi u_i C_{il}^\phi S_{il} \right)
\end{aligned}$$

where  $C_i^\phi$  and  $C_{il}^r$  are diagonal matrices with confidence parameters  $c_{ij}^\phi$  and  $c_{ijl}^r$ . We define  $S$  as  $N \times N$  matrix,  $\phi$  as  $K \times N^2$  matrix and  $R_j$  as vector with  $r_{ijl}$  values for all pairs of users  $i$  and  $l$  for the given item  $j$ . Given updated variables  $\{u_i, v_j, \phi_{il}, s_i\}$  the topic proportions  $\theta_j$  is updated by applying Jensen's inequality (Wang and Blei 2011):

$$\begin{aligned}
\ell(\theta_j) &\geq \sum_m^{W(j)} \sum_k^K \psi_{jmk} (\log \theta_{jk} \beta_{k, w_{jm}} - \log \psi_{jmk}) \\
&\quad - \frac{\lambda_v}{2} \sum_j^D (v_j - \theta_j)^T (v_j - \theta_j) \\
&= \ell(\theta_j, \psi_j)
\end{aligned} \tag{4}$$

where  $\psi_{jmk} = q(z_{jm} = k)$  and  $\ell(\theta_j, \psi_j)$  gives the tight lower bound of  $\ell(\theta_j)$ . The optimal  $\psi_{jmk}$  satisfies  $\psi_{jmk} \propto \theta_{jk} \beta_{k, w_{jm}}$ . Given updated variables  $\{u_i, v_j, \phi_{il}, s_i, \theta_j\}$ , we can optimize  $\beta$ ,

$$\beta_{kw} \propto \sum_j \sum_m \psi_{jmk} \delta(w_{jm} = w) \tag{5}$$

where  $\delta$  is one if and only if  $w$  is assigned to  $w_{jm}$ .

## Prediction

After all the parameters are learned,  $\mathcal{L}A$ -CTR model can be used for both in-matrix and out-of-matrix prediction with either user's attention ( $\phi$ ) or interest ( $u$ ). As (Wang and Blei 2011) mentioned, in-matrix prediction refers to the prediction task for user's rating on item that has been rated at least once by other users, while out-of-matrix prediction refers to predicting user's rating on a new item that has no rating history.

For in-matrix prediction with attention ( $\phi$ ), the prediction that user  $i$ 's ratings for item  $j$  via user  $l$  is obtained by point estimation with optimal variables ( $\phi^*$ ,  $u^*$ ,  $v^*$ ,  $\theta^*$ ):

$$\begin{aligned}
\mathbb{E}[r_{ijl} | \mathcal{D}] &\approx \mathbb{E}[\phi_{il} | \mathcal{D}]^T (\mathbb{E}[\theta_j | \mathcal{D}] + \mathbb{E}[\epsilon_j | \mathcal{D}]) \\
r_{ij}^* &\approx \phi_{il}^{*T} v_j^*
\end{aligned} \tag{6}$$

where the rating of user  $i$  is decided by user's attention  $\phi_{il}^*$  and item topic profile  $v_j^*$ .

For out-of-matrix prediction with attention ( $\phi$ ), when the item is new and no ratings are observed, the point estimation equation is:

$$\begin{aligned}
\mathbb{E}[r_{ijl} | \mathcal{D}] &\approx \mathbb{E}[\phi_{il} | \mathcal{D}]^T \mathbb{E}[\theta_j | \mathcal{D}] \\
r_{ij}^* &\approx \phi_{il}^{*T} \theta_j^*
\end{aligned} \tag{7}$$

where item topic profile is  $\theta_j^*$  due to  $\epsilon_j^* = 0$ .

For in-matrix prediction with user interest ( $u$ ), the point estimation prediction that user  $i$ 's rating for item  $j$  is:

$$\begin{aligned}
\mathbb{E}[r_{ij} | \mathcal{D}] &\approx \mathbb{E}[u_i | \mathcal{D}]^T (\mathbb{E}[\theta_j | \mathcal{D}] + \mathbb{E}[\epsilon_j | \mathcal{D}]) \\
r_{ij}^* &\approx u_i^{*T} v_j^*
\end{aligned} \tag{8}$$

where the prediction of user's rating is decided by user interest profile  $u_i^*$  and item topic profile  $v_j^*$ .

For out-of-matrix prediction with user interest ( $u$ ):

$$\begin{aligned}
\mathbb{E}[r_{ij} | \mathcal{D}] &\approx \mathbb{E}[u_i | \mathcal{D}]^T \mathbb{E}[\theta_j | \mathcal{D}] \\
r_{ij}^* &\approx u_i^{*T} \theta_j^*
\end{aligned} \tag{9}$$

where the rating of a user is decided by the user's interest profile  $u_i^*$  and item's topic profile  $\theta_j^*$ .

## Evaluation

In this section we demonstrate the utility of the  $\mathcal{L}A$ -CTR model both with interest ( $u$ ) and attention ( $\phi$ ) hidden variables. We apply the  $\mathcal{L}A$ -CTR model to data from the social news aggregator Digg. We compare our results to state-of-the-art alternative models: CTR and CTR-smf.

## Data Set

Digg allows users to submit links to news stories and other users to recommend stories by voting for or "digging" stories they find interesting. Digg also allows users to follow the activity of other users to see the stories they submitted or voted for recently. When a user votes for a story, all her followers can see the vote, so the information is broadcast and shared with all social network neighbors. At the time datasets were collected, users were submitting tens of thousands of stories, from which Digg selected a handful (about 100) to promote to its front page based on how popular in the community.

We perform evaluations on two datasets: the 2009 dataset (Lerman and Ghosh 2010) and the 2010 dataset (Sharara, Rand, and Getoor 2011). The 2009 dataset contains information about the voting history of 70K active users on 3.5K stories promoted to Digg front page in June 2009, and contains 2.1 million votes. The follower graph of contains 1.7 million social links. At the time this dataset was collected, Digg was assigning stories to one of 8 topics (Entertainment, Lifestyle, Science, Technology, World & Business, Sports, Offbeat, and Gaming) and one of 50 subtopics (World News, Tech Industry News, General Sciences, Odd Stuff, Movies, Business, World News, Politics, etc.). The 2010 data set contains information about voting histories of 12K users over a 6 months period (Jul - Dec 2010). It includes 48K stories with 1.9 million votes and contains 1.3 million social links in the follower graph. At the time data was collected, Digg was assigning stories to 10 topics, replacing the 2009 "World & Business" category with "World News," "Business," and "Politics" and dropping lower-level topic categories.

We examine only the votes that the story accrues before promotion to the front page. During that time, the story propagates mainly via friends' votes (or recommendations), although some users could discover the story on the Upcoming

stories page, which received tens of thousands of submissions daily. After promotion, users are likely to be exposed to the story through the front page, and vote for it independently of friends’ recommendations. In the 2009 data set, 28K users voted for 3,553 stories and in the 2010 data set, 4K users voted for 36,883 stories before promotion. We use tf-idf to choose the top 3K distinct words from the titles and user-supplied summaries of stories in the 2009 data set and 4K distinct words in the 2010 data set as the vocabularies. We focused the data further by selecting only those users who voted at least 10 times and items containing more than 10 words, resulting in 1.4K users (who voted for 3K stories) in the 2009 data set and 1.8K users (who voted on 18K stories) in the 2010 data set.

### Evaluation on Vote Prediction

We consider a vote by user  $i$  on story  $j$  as an item adoption event, or a rating of one. A rating of zero means either that the user did not like the item or was not aware of it; therefore, we only consider rated items for evaluation. For a fair and easy comparison with CTR and CTR-smf, we use  $recall@X$ . We present each user with  $X$  items sorted by their predicted rating score, and evaluate based on the fraction of items that the user actually voted for. We define  $recall@X$  for a user as:

$$recall@X = \frac{\text{num of items in top } X \text{ user votes for}}{\text{total num of items user votes for}} \quad (10)$$

and we average recall values of all users to summarize performance of the algorithm. Note a better model should provide higher  $recall@X$  at each  $X$ .

### Impact of Parameters




Figure 2: Average  $recall@100$  of in-matrix prediction for Digg 2009 by varying precision parameters  $\lambda_\phi$  while other parameters are fixed.

We study how parameters of  $\mathcal{L}A$ -CTR affect the overall performance of in-matrix prediction using  $recall@100$ . The impact of parameter  $\lambda_v$  has been studied by CTR and CTR-smf papers, and we found similar trends. Here, we show how the attention parameter ( $\lambda_\phi$ ) affects the overall performance of  $\mathcal{L}A$ -CTR. We vary  $\lambda_\phi \in \{0, 0.001, \dots, 1000\}$ , while we fix other parameters:  $\lambda_v = 100$ ,  $\lambda_u = 0.01$ , and  $\lambda_s = 0.01$ . A larger value of  $\lambda_\phi$  increases the penalty of user’s attention ( $\phi$ ) diverging from interest ( $u$ ) and influence ( $s$ ). In Figure 2,

we show the impact of  $\lambda_\phi$  on  $\mathcal{L}A$ -CTR performance with only user interests ( $u$ ) and with attention ( $\phi$ ). With small values of  $\lambda_\phi$ , user’s interest ( $u$ ) and influence ( $s$ ) highly diverge from his attention ( $\phi$ ). It explains why  $recall@100$  with small  $\lambda_\phi$  provides poor performance on recommendation tasks with user interests ( $u$ ) only. Similarly large values of  $\lambda_\phi$  will increase the penalty of user’s attention diverging from user’s influence and interests. Best performance is achieved when  $\lambda_\phi = 1.0$ , for both recommendations with user interest ( $u$ ) and user attention ( $\phi$ ).

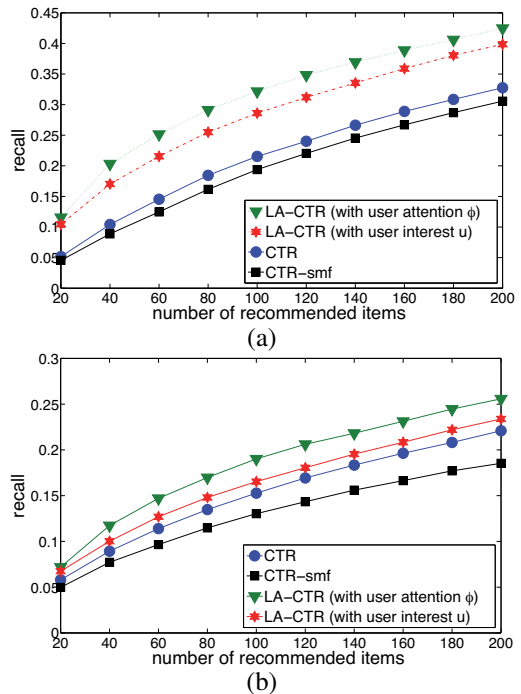


Figure 3: Recall of in-matrix prediction for (a) Digg 2009 and (b) 2010 data set by varying the number of recommended items ( $@X$ ) with the 200 top tips.

### Model Selection

For collaborative topic regression (CTR), we set the parameters  $K=200$ ,  $\lambda_u = 0.01$ ,  $\lambda_v = 100$ ,  $a = 1$ ,  $b = 0.01$  by using grid search on held out recommendations. The precision parameters  $\lambda_v$  balances how the item’s latent vector  $v_j$  diverges from the topic proportions  $\theta_j$ . We vary  $\lambda_v \in \{0.001, 0.01, 0.1, 1, 10, 100, 1000\}$ , where a larger  $\lambda_v$  increases the penalty of  $v_j$  diverging from  $\theta_j$ . For Collaborative Topic Regression with social matrix factorization (CTR-smf), we choose the parameters using grid search on held out recommendations. Through out this paper, we set parameters  $K=200$ ,  $\lambda_u = 0.01$ ,  $\lambda_v = 100$ ,  $\lambda_q = 10$ ,  $\lambda_s = 2$ ,  $a=1$ , and  $b=0.01$ .

For  $\mathcal{L}A$ -CTR, we choose the parameters similar to CTR and CTR-smf approach by using grid search on held out recommendations. We set parameters  $K=200$ ,  $\lambda_u = 0.01$ ,  $\lambda_v = 100$ ,  $\lambda_\phi = 1$ ,  $\lambda_s = 0.01$ ,  $a^\phi=1$ ,  $a^r=1$ ,  $b^\phi=0.01$ , and  $b^r=0.01$ .

## Vote Prediction

Next, we evaluate the proposed  $\mathcal{L}A$ -CTR model by measuring how well it allows us to predict individual votes. There are 257K pre-promotion votes in the 2009 dataset and 1.5 million votes in the 2010 dataset, with 72.34 and 68.20 average votes per story, respectively. For our evaluation, we randomly split the data into training and test sets, and performed five-fold cross validation. The average percentage of positive examples in the test set is 0.73%, suggesting that friends share many stories that users end up not voting for. High sparsity of dataset makes the prediction task extremely challenging, with less than one in a hundred chance of successfully predicting votes if stories are picked randomly.

We train the models on the data with selected models of CTR, CTR-smf,  $\mathcal{L}A$ -CTR and perform in-matrix prediction while we vary the number of top  $X$  items  $\{20, 40, \dots, 200\}$ . Figure 3 shows the  $recall@X$  for in-matrix prediction for CTR, CTR-smf,  $\mathcal{L}A$ -CTR both with interest and attention. All four models show performance improvement when the number of returned items is increased. Our model with attention always outperformed other three models consistently. Even with learned user interests,  $\mathcal{L}A$ -CTR model always outperformed CTR and CTR-smf. Note that recall scores with smaller than 100 of  $\mathcal{L}A$ -CTR with user attention are already higher than ones of CTR-smf  $recall@200$ .

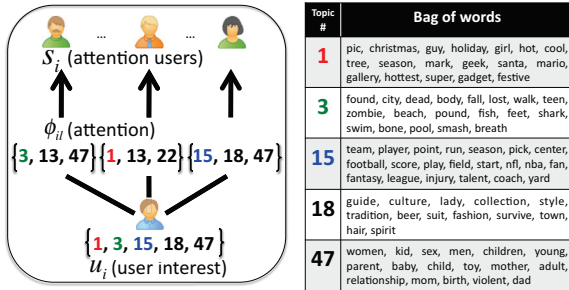


Figure 4: One example user from Digg 2009 data set. We show top 5 topics in learned  $u_i$  and top 3 topics in learned attention  $\phi_{il}$  to his three most influential users. We also show the bag of words that represents each topic.

## Limited Attention

An important advantage of  $\mathcal{L}A$ -CTR is that it can explain both course-grained and fine-grained levels of latent topic spaces: (1) user interest topic space ( $u_i$ ) and (2) attention topic space ( $\phi_{il}$ ) using the topics learned from the corpus. Furthermore, it also explains the influence of others ( $s_l$ ) on user's item adoption decisions. In Figure 4, we show one example from the 2009 data set. With  $\mathcal{L}A$ -CTR we can find the top  $n$  matched topics by ranking the elements of his interest vector  $u_i$ . In this example, user  $i$  is interested in topics 1, 3, 15, 18, and 47. Furthermore, we can also explain influence weight of others on user  $i$  as well as user  $i$ 's attention on user  $l$  over topic space. Here, user  $i$  pays attention to the last user mostly on sports topic (topic # 15), while he pays attention the first two users on other topics.




Figure 5: Recall of in-matrix prediction for Digg 2009 by varying the number of recommended items ( $@X$ ) with the 50 topics.

## Conclusion

In this paper, we proposed  $\mathcal{L}A$ -CTR for recommending items to users in online social media based on users' limited attention as well as their ratings, their social network, and the content of items. We showed that by taking into account users' attention, the proposed model outperforms other state-of-art algorithms on item adoption prediction task. Our model provides not only interpretable user profiles, but also fine-grained description of how much attention the user pays to others on which topics. Such description could help explain why people follow others and how information propagates over the online social network.

One disadvantage of  $\mathcal{L}A$ -CTR is its model complexity, since the size of user-item rating is  $N^2D$ . However, due to high sparsity of the social network and the user-item matrix we did not experience any major slow down. Furthermore, given that with the small number of topics,  $\mathcal{L}A$ -CTR already outperformed the other models. In Figure 5, we show recall of in-matrix prediction for Digg 2009 data set with 50 topics. Not only does  $\mathcal{L}A$ -CTR model outperform both CTR and CTR-smf with 50 topics, but it also outperformed both CTR and CTR-smf with 200 topics. In other words, increasing the size of the topic space and descriptive power of competitor models did not lead to better performance compared to  $\mathcal{L}A$ -CTR. In the future, we will optimize  $\mathcal{L}A$ -CTR speed by applying fast bayesian posterior sampling (i.e. stochastic gradient fisher scoring) and will apply it to larger data sets.

For future work, we are interested in applying our model to online social media data set with actual source of adoption information. We will test whether we are able to learn users limited attention and their interest latent vector better when the actual source of adoption information is available. Note that, our model not only predicts the most likely items but also predicts the most likely source of adoption. With the available source of adoption data set, we will also test our source of adoption predictions for the new item adoptions.

## Acknowledgment

This material is based upon work supported by the Air Force Office of Scientific Research under Contract Nos. FA9550-10-1-0102 and FA9550-10-1-0569, and by DARPA under Contract No. W911NF-12-1-0034.

## References

- Agarwal, D., and Chen, B.-C. 2009. Regression-based latent factor models. In *KDD*, 19–28.
- Agarwal, D., and Chen, B.-C. 2010. fLDA: Matrix factorization through latent dirichlet allocation. In *WSDM*.
- Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone’s an influencer: quantifying influence on twitter. In *WSDM*.
- Blei, D.; Ng, A.; and Jordan, M. 2003. Latent dirichlet allocation. *The Journal of Machine Learning Research* 3:993–1022.
- Chang, J. 2009. Relational topic models for document networks. In *AISTATS*.
- Chua, F. C. T.; Lauw, H. W.; and Lim, E.-P. 2012. Generative models for item adoptions using social correlation. *IEEE Transactions on Knowledge and Data Engineering* 99.
- Counts, S., and Fisher, K. 2011. Taking it all in? visual attention in microblog consumption. In *ICWSM*.
- Daumé, III, H. 2009. Markov random topic fields. In *Short Papers at ACL-IJCNLP*, 293–296.
- Gilbert, E., and Karahalios, K. 2009. Predicting tie strength with social media. In *CHI*.
- Goncalves, B.; Perra, N.; and Vespignani, A. 2011. Validation of Dunbar’s number in Twitter conversations. *arXiv.org*.
- Herlocker, J.; Konstan, J.; Borchers, A.; and Riedl, J. 1999. An algorithmic framework for performing collaborative filtering. In *ACM SIGIR conference on Research and development in information retrieval*, 230–237. ACM.
- Hodas, N., and Lerman, K. 2012. How limited visibility and divided attention constrain social contagion. In *SocialCom*.
- Kahneman, D. 1973. *Attention and effort*. Prentice Hall.
- Kang, J.-H.; Lerman, K.; and Getoor, L. 2013. La-lda: A limited attention model for social recommendation. In *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction (SBP)*.
- Karypis, G. 2000. Evaluation of item-based top-n recommendation algorithms. Technical report, DTIC Document.
- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37.
- Lerman, K., and Ghosh, R. 2010. Information contagion: an empirical study of spread of news on digg and twitter social networks. In *ICWSM*.
- Ma, H.; Yang, H.; Lyu, M. R.; and King, I. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and knowledge management*, 931–940. ACM.
- Nallapati, R., and Cohen, W. 2008. Link-PLSA-LDA: A new supervised model for topics and influence of blogs. In *ICWSM*.
- Purushotham, S.; Liu, Y.; and Kuo, C. 2012. Collaborative topic regression with social matrix factorization for recommendation systems. *arXiv preprint arXiv:1206.4684*.
- Ramage, D.; Hall, D.; Nallapati, R.; and Manning, C. 2009. Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora. In *EMNLP*.
- Rensink, R.; O’Regan, J.; and Clark, J. 1997. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 8(5):368.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In *WWW*.
- Rosen-Zvi, M.; Griffiths, T.; Steyvers, M.; and Smyth, P. 2004. The author-topic model for authors and documents. In *UAI*.
- Salakhutdinov, R., and Mnih, A. 2008. Probabilistic matrix factorization. *Advances in neural information processing systems* 20:1257–1264.
- Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2001. Itembased collaborative filtering recommendation algorithms. In *WWW*.
- Sharara, H.; Rand, W.; and Getoor, L. 2011. Differential adaptive diffusion: Understanding diversity and learning whom to trust in viral marketing. In *ICWSM*.
- Wang, C., and Blei, D. M. 2011. Collaborative topic modeling for recommending scientific articles. In *KDD*.
- Weng, L.; Flammini, A.; Vespignani, A.; and Menczer, F. 2012. Competition among memes in a world with limited attention. *Scientific Reports* 2.
- Yang, S.-H.; Long, B.; Smola, A.; Sadagopan, N.; Zheng, Z.; and Zha, H. 2011. Like like alike: joint friendship and interest propagation in social networks. In *Proceedings of the 20th international conference on World wide web*, 537–546. ACM.
- Yu, K.; Lafferty, J.; Zhu, S.; and Gong, Y. 2009. Large-scale collaborative prediction using a nonparametric random effects model. In *ICML*, 1185–1192.