# Latent Space Model for Multi-Modal Social Data

Yoon-Sik Cho<sup>†</sup>, Greg Ver Steeg<sup>‡</sup>, Emilio Ferrara<sup>‡</sup>, Aram Galstyan<sup>‡</sup>

University of Southern California, Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA 90292 USA <sup>†</sup>yoonsik.cho@usc.edu, <sup>‡</sup>{gregv, ferrarae, galstyan }@isi.edu

# ABSTRACT

With the emergence of social networking services, researchers enjoy the increasing availability of large-scale heterogenous datasets capturing online user interactions and behaviors. Traditional analysis of techno-social systems data has focused mainly on describing either the dynamics of social interactions, or the attributes and behaviors of the users. However, overwhelming empirical evidence suggests that the two dimensions affect one another, and therefore they should be jointly modeled and analyzed in a multi-modal framework. The benefits of such an approach include the ability to build better predictive models, leveraging social network information as well as user behavioral signals. To this purpose, here we propose the Constrained Latent Space Model (CLSM), a generalized framework that combines Mixed Membership Stochastic Blockmodels (MMSB) and Latent Dirichlet Allocation (LDA) incorporating a constraint that forces the latent space to concurrently describe the multiple data modalities. We derive an efficient inference algorithm based on Variational Expectation Maximization that has a computational cost linear in the size of the network, thus making it feasible to analyze massive social datasets. We validate the proposed framework on two problems: prediction of social interactions from user attributes and behaviors, and behavior prediction exploiting network information. We perform experiments with a variety of multi-modal social systems, spanning location-based social networks (Gowalla), social media services (Instagram, Orkut), e-commerce and review sites (Amazon, Ciao), and finally citation networks (Cora). The results indicate significant improvement in prediction accuracy over state of the art methods, and demonstrate the flexibility of the proposed approach for addressing a variety of different learning problems commonly occurring with multi-modal social data.

# Keywords

multi-modal social networks; topic models; LDA

# 1. INTRODUCTION

The proliferation of massive-scale social networking services produces extensive amounts of data describing various forms of user behaviors along with detailed information about social interactions. Traditionally, studies of online social systems have focused on analyzing the networks induced by such interactions [1, 42, 15, 30, 38, 43, 54, 21, 19] and examining the impact of structural features on a variety of dynamical processes unfolding on the network (e.g., information diffusion, cascades, etc.) [10, 39, 11, 26, 5, 6, 23, 44, 13, 24]. However, a number of recent studies makes it increasingly clear that focusing only on network properties while discarding rich contextual information limits our comprehensive understanding of complex social systems dynamics [52, 36, 4, 56, 18, 32, 49, 37, 20, 25, 27, 2].

In addition to social network data, we often have information about user behaviors such as mobility patterns (e.g., Foursquare check-ins), user generated content (e.g., tweets or Facebook posts), purchase history, and so on. User behaviors and social interactions are often closely linked. This relationship has been extensively studied in social sciences under the umbrella of *homophily theory* [41], which states that users exhibiting similar attributes or activity patterns are also more likely to be socially linked. When user attributes relevant to forming social links are not directly observable, this phenomenon is called *latent homophily*.

From the predictive modeling perspective, homophily (or its opposite, *heterophily*) can be used to build more accurate models of user behavior and social interactions based on multi-modal data. Such an approach can generate a more comprehensive understanding of users and their preferences [57, 48, 46]. Furthermore, these types of models can be employed to generate predictions across modalities. For instance, we can try forecast the social interactions of an individual based on (partial) knowledge of her/his behaviors. Conversely, we can predict user behavior based on available social network information. This type of capability is particularly important where we have sparse data in one single modality, yet data across multiple modalities abound. Thus, there is an overwhelming need for developing a unified computational framework to describe and analyze multi-modal social data.

Latent space models [29] provide a viable framework to analyze multi-modal social data. Indeed, models based on *Latent Dirichlet Allocation* (LDA) [9] and its extensions have been recently used to analyze user-generated content [47], mobility patterns [34], etc. At the same time, latent space models such as *Mixed Membership Stochastic Blockmodels* 

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media. *WWW 2016*, April 11–15, 2016, Montréal, Québec, Canada. ACM 978-1-4503-4143-1/16/04. http://dx.doi.org/10.1145/2872427.2883031.

(MMSB) have been devised to describe multi-faceted interactions in social networks [3].

Using a latent space model that jointly represents information about user behaviors (or attributes) and social interactions is an attractive approach since it provides an explicit and concrete mechanism to explain the observed correlations between the two. Pairwise Link-LDA [45] combines LDA [9] and MMSB [3] to describe both the network structure and user-generated content. Unfortunately, Pairwise Link-LDA often fails to exploit the synergies between different modalities, as it tends to *fragment* the latent space into non-overlapping (or weakly overlapping) regions corresponding exclusively to either user attributes or to the social network. This fragmentation is due to the weak dependency between the attributes and links which are connected only through a common latent topic distribution. To address this shortcoming, the *Relational Topic Model* (RTM) [12] was designed to impose additional constraints on the generative model that would require stronger correlations between links and attributes. However, RTM utilizes a significantly simplified model of link formation that does not account for truly multi-faceted interactions among the users. A detailed technical account of these existing methods and their limitations is provided in Section  $\S2.1$ .

In this paper, we propose an alternative approach to modeling multi-modal social data that combines the benefits of Pairwise Link-LDA and RTM while avoiding their shortcomings. Our proposal, called Constrained Latent Space Model (CLSM), augments the generative process by introducing a set of constraints that allows to account for stronger correlations between user attributes and social networks, and avoid the fragmentation of the joint latent space. In contrast to other methods such as RTM, CLSM uses the full generative process of MMSB to describe the link formation, which yields a truly multi-faceted modeling of social interactions. Our experiments in link and behavior prediction reveal that the latter feature is particularly important when nodes can interact with each other in multiple modes/dimensions. Overall, the proposed method achieves significant improvement in performance over the others in various learning scenarios.

Our primary contributions can be summarized as follows:

- We propose a general framework called Constrained Latent Space Model (CLSM) that combines MMSB and LDA with an added constraint that forces the latent space to jointly describe multiple data modalities.
- We derive an efficient inference and learning algorithm based on *Variational EM* (expectation-maximization), which makes use of the constraint to significantly reduce the computational time to a linear cost in the size of the network, allowing for large-scale data modeling.
- We perform extensive experiments on datasets that augment social network information with various data modalities including texts, purchase reviews, and checkin data. Our framework improves the state of the art both in link prediction and in user behavior prediction.
- We analyze two case studies in detail, the citation network *Cora* and the location-based social network *Gowalla*, discussing the unique insights CLSM yields over using other multi-modal latent space models.

The rest of the paper is organized as follows: Sec. §2 introduces the background on latent space models and mentions methods related to our proposal. Sec. §3 presents the Constrained Latent Space Model (CLSM), and describes the Variational Expectation Maximization algorithm. A rigorous experimental evaluation for link and attribute prediction is provided in Sec. §4, with benchmarks on synthetic data and analyses of six real-world multi-modal social systems. Sec. §5 summarizes our contribution and future plans.

# 2. BACKGROUND

Consider a set of N nodes that form a network Y:  $y(n_1, n_2) = 1$  if nodes  $n_1$  and  $n_2$  are linked, and  $y(n_1, n_2) = 0$  otherwise. While the model considered here is quite general, we will focus on undirected networks where  $y(n_1, n_2) = y(n_2, n_1)$ .

In addition to network information, we also have information about node behavior, such as purchasing a certain item, or visiting a certain location. We assume that the behaviors are selected from a finite set, or vocabulary,  $\mathcal{V} = \{1, ..., V\}$ . We characterize the selection  $v \in \mathcal{V}$  by a vector w of size V, so that all the components of w are zero except the v-th component, which is equal to  $w^v = 1$ . Let  $M_n$  be the number of times we have observed node n's selections. Then the nodes's behavior is represented by  $\mathbf{w}_n = (\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_{M_n})$ .

We assume that the behavior  $\mathbf{w}_n$  of user n is characterized by a latent *topic* distribution  $\boldsymbol{\theta}_n$  defined over a Kdimensional simplex, where K is the number of topics. In the context of the present work, a topic can be interpreted as a behavior type. Each topic itself is a distribution over Vpossible selections (locations, purchases, etc.), denoted here by  $\boldsymbol{\omega}_k, \ k = 1, \ldots, K$ . In the LDA model used here, the topics are sampled from a Dirichlet prior,  $\boldsymbol{\theta}_n \sim \text{Dir}(\alpha)$ . To generate the m-th component of the attribute vector for user n, we first sample a topic indicator variable from a multinomial distribution,  $\mathbf{z}_{n,m} \sim \text{Mult}(\boldsymbol{\theta}_n)^1$ , and then sample a selection  $\mathbf{w}_{n,m}$  given the selected topic and its corresponding multinomial distribution:  $\mathbf{w}_{n,m} \sim \text{Mult}(\boldsymbol{\omega}_i)$ , where i is the index of the only non-zero component in  $\mathbf{z}_{n,m}$ .

The main premise behind the joint latent space modeling approach is that the same topic distribution also determines the node's social interactions, where the topics now can be interpreted as different groups or communities a node is affiliated with. Namely, for a pair of nodes  $(n_1, n_2)$ , we sample two indicator variables  $\mathbf{z}_{n_1 \to n_2}$  and  $\mathbf{z}_{n_2 \leftarrow n_1}$ , and then establish a link between those two nodes with probability  $f(\mathbf{z}_{n_1 \to n_2}, \mathbf{z}_{n_1 \leftarrow n_2})$ , where f is a link function. For instance, in Mixed Membership Stochastic Blockmodels, f is parameterized via a  $K \times K$  compatibility matrix B so that  $B_{k_1k_2}$ is the probability of establishing a link between nodes from groups  $k_1$  and  $k_2$ . When the off-diagonal components in B are zero, the only allowed interactions are intra-group ones, which yields strictly assortative Mixed Membership Stochastic Blockmodels (aMMSB) [28]. We distinguish aMMSB from conventional MMSB [3] by naming the latter model as full MMSB (fMMSB), to reflect the use of a full compatibility matrix. Consistent with previous reports [28], we found that aMMSB outperforms fMMSB for both Pairwise Link-LDA and our proposed model. Hence, below we always use aMMSB for network generation.

 $<sup>{}^{1}\</sup>mathbf{z}_{n,m}$  is a size-K vector where all the components are zero except the component corresponding to the multinomial selection, which equals to 1.

## 2.1 Related Work

Before describing the proposed model, we review two related approaches that are the closest to our model, Pairwise Link-LDA [45] and RTM [12].

Both approaches use an LDA-based process for generating behaviors. Specifically, for each user n, we sample topic weights  $\boldsymbol{\theta}_n$  from a Dirichlet prior  $\boldsymbol{\alpha}$ , and a number of selections  $M_n$  from a Poisson distribution. Then, for each selection m, we sample a topic indicator  $\mathbf{z}_{n,m}$ , and generate attributes from the corresponding multinomial distribution,  $\mathbf{w}_{n,m} \sim \text{Mult}(\boldsymbol{\omega}_{\mathbf{z}_{n,m}}).$ 

The two approaches differ in their link formation model. In Pairwise Link-LDA, the links are generated according to the Mixed Membership Stochastic Blockmodel [3]. Namely, given two nodes  $n_1$  and  $n_2$ , we sample indicator variables from the corresponding topic distributions,  $\mathbf{z}_{n_1 \to n_2} \sim \text{Mult}(\boldsymbol{\theta}_{n_1})$ ,  $\mathbf{z}_{n_1 \leftarrow n_2} \sim \text{Mult}(\boldsymbol{\theta}_{n_2})$ , and then generate a link according to a Bernoulli trial,  $y(n_1, n_2) \sim \text{Bernoulli}(\mathbf{z}_{n_1 \to n_2}^\top \text{B}\mathbf{z}_{n_1 \leftarrow n_2})$ , where B is the  $K \times K$  compatibility matrix.

From the joint modeling perspective, Pairwise Link-LDA imposes correlations between user behavior and social interactions by sampling the corresponding indicator variables from the same (multinomial) distributions. As we mentioned above, this framework fails to properly account for observed correlations, especially when the number of topics is large [12]. Instead, the latent space is fragmented into network-specific and behavior-specific subspaces. A similar phenomenon was observed in [8] where the authors used a shared latent space for describing images and their captions.

The Relational Topic Model (RTM) [12] tried to address this problem by modifying the link formation mechanism. Namely, instead of sampling new indicator variables for link formation, RTM reuses the indicator variables used for generating behaviors. Let  $\bar{\mathbf{z}}_n$  be the mean of the indicator variables used for generating behaviors,  $\bar{\mathbf{z}}_n = \frac{1}{M_n} \sum \mathbf{z}_{n,m}$ . Then, given a pair of nodes  $(n_1, n_2)$ , a link is sampled via  $y(n_1, n_2) \sim \text{Bernoulli}(f_{RTM}(\bar{\mathbf{z}}_{n_1} \circ \bar{\mathbf{z}}_{n_2}))$ , where  $\circ$  denotes the Hadamard (element-wise) product, and  $f_{RTM}(\cdot)$  is a parameterized link function, e.g., sigmoidal or exponential [12].

While RTM is superior to Pairwise Link-LDA for describing correlations between user attributes and interactions, its simplified link formation mechanism does not have the flexibility of MMSB, and is not adequate for capturing truly heterogenous and multi-faceted interactions between the users. Indeed, as illustrated in Section 4.1, RTM is not able to accurately model scenarios where the nodes have well-mixed memberships. We next formulate our proposed generative model that addresses the shortcomings of both methods.

# 3. CONSTRAINED LATENT SPACE MODEL

This section describes our proposal, the *Constrained La*tent Space Model (CLSM). The plate diagram of the generative process is illustrated in Figure 1, which for sake of comparison also depicts the generative process for Pairwise Link-LDA. Both approaches use the same scheme to generate the social network. Where the approaches differ, however, is the generative process for the user behaviors (attributes). Namely, when generating behaviors for user n, we reuse the indicator variables that were used to generate links for the same user. In this regard, our approach shares some similarity with RTM, in that it also uses the same set of indicator variables for both attribute and link generation. In contrast to RTM, however, we first generate a set of indicator variables by sampling the links, and then use this set to generate attributes. Remarkably, our experiments show that the order of sampling is not important, and we obtain the same results by first sampling attributes, as done in RTM. However, sampling the links first is computationally more efficient when the network is relatively sparse, which is the case for the vast majority of datasets reflecting realworld techno-social systems. In the opposite case of dense networks and sparse attributes, one can sample the other way around for efficiency.

Following the language of MMSB, let N be the total number of nodes and K be the total number of network communities (that is, topics). The generative process is as follows:

#### • Network Generation:

- 1. For each community k, sample the community strength  $\beta_k \sim \text{Beta}(\eta_1, \eta_0)$ , where  $\beta_k := \mathsf{B}(k, k)$
- 2. For each node n, sample the  $K \times 1$  membership vector  $\boldsymbol{\theta}_n \sim \text{Dirichlet}(\boldsymbol{\alpha})$ .
- 3. For each node n, initialize an empty multiset of indicator variables,  $\mathbb{Z}_n = \emptyset$ .
- 4. For each pair  $(n_1, n_2)$ 
  - (a) Draw membership indicator vectors  $\mathbf{z}_{n_1 \to n_2} \sim \operatorname{Mult}(\boldsymbol{\theta}_{n_1})$  $\mathbf{z}_{n_1 \leftarrow n_2} \sim \operatorname{Mult}(\boldsymbol{\theta}_{n_2}).$
  - (b) Sample the pair interaction  $y(n_1, n_2) \sim \text{Bernoulli}(\mathbf{z}_{n_1 \to n_2}^{\top} \mathsf{B} \mathbf{z}_{n_1 \leftarrow n_2}),$ where  $\mathsf{B} = \text{diag}(\beta_1, \beta_2, ..., \beta_K) + \epsilon(\mathbb{1} - I),$ and  $\epsilon$  is a small regularization parameter.
  - (c) Augment the corresponding indicator multisets (allowing repetitions).  $\mathbb{Z}_{n_1} \to \mathbb{Z}_{n_1} \cup \{\mathbf{z}_{n_1 \to n_2}\}$  $\mathbb{Z}_{n_2} \to \mathbb{Z}_{n_2} \cup \{\mathbf{z}_{n_1 \leftarrow n_2}\}$

#### • Behavior Generation:

Let  $M_n$  be the total number of selections<sup>2</sup> of user n from a (behavior) set  $\mathcal{V} = \{1, ..., V\}$ .

- 1. For each topic k, sample the attribute distribution  $\omega_k \sim \text{Dirichlet}(\kappa)$ , where  $\kappa$  is a hyperparameter.
- 2. For each selection  $\mathbf{w}_{n,m}$  of user n, where  $m \in \{1, ..., M_n\}$ ,

  - (b) Sample a selection  $\mathbf{w}_{n,m} \sim \text{Mult}(\boldsymbol{\omega}_i)$ , where i is the index of  $\hat{\mathbf{z}}$ 's non-zero component.

Exact inference with the above generative model is not feasible due to the presence of latent variables [29, 7]. In the following, we describe an approximate method for learning and inference based on Variational Expectation Maximization (EM). We will use it to estimate the parameters  $\{\beta\}, \{\omega\}$  and infer the posterior distributions of the model's latent variables. The computational cost of the this inference method is linear in the size of the network (see section §3.1), which makes CLSM a feasible tool for analyzing large-scale multi-modal social data.

<sup>&</sup>lt;sup>2</sup>For the purposes of data generation  $M_n$  can be sampled from, say, a Poisson distribution. This is not relevant for inference, however, where  $M_n$  is specified in the data.



Figure 1: Plate Diagram comparison of Pairwise Link-LDA (a) and CLSM (b). Shaded nodes represent the observable variable. (a) Pairwise Link-LDA focusing on a pair: attributes and links are generated from the same latent variable, but a dependency between the latent factors  $\mathbf{z}_{n\to n'}$  and  $\mathbf{z}_n$  lacks. (b) CLSM: we focus on node n and the behaviors of other nodes are omitted. Dot represents all possible nodes interacting with node n.  $\mathbf{z}_n$  now has high dependency on the set of  $\mathbf{z}_{n\to ..}$ .

#### 3.1 Variational Inference

Given the observation of the interaction network  $\mathbf{Y}$  and user behaviors  $\mathbf{w}_{1:N} = {\{\mathbf{w}_k\}_{k=1}^N}$ , we are interested in inferring the posterior distributions of the model's latent variables,  $p(\boldsymbol{\theta}_{1:N}, \mathbb{Z}_{1:N}, \mathbf{C} | \mathbf{Y}, \mathbf{w}_{1:N})$  (where  $\mathbf{C}$  is the collection of all  $\mathbf{c}_m$  of all nodes), as well as estimating the hyperparameters  $\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\kappa}$ . A number of approximate inference algorithms have been proposed in literature. In this paper, we rely on variational method [33] that approximates the posterior by a computationally tractable variational distribution with some free variational parameters. Those parameters are selected to minimize the Kullback-Leibler divergence between the variational distribution and the true posterior.

Here we suggest a factorized variational distribution over the latent variables  $q(\boldsymbol{\theta}_{1:N}, \mathbb{Z}_{1:N}, \mathbf{C})$ :

$$q(\boldsymbol{\theta}_{1:N}, \mathbb{Z}_{1:N}, \mathbf{C}) = \prod_{n=1}^{N} q_{\text{dir}}(\boldsymbol{\theta}_{n} | \boldsymbol{\gamma}_{n}) \prod_{m=1}^{M_{n}} q_{\text{mul}}(\mathbf{c}_{m}^{n} | \boldsymbol{\lambda}_{m}^{n})$$
$$\prod_{n_{1}, n_{2}}^{N} q_{\text{mul}}(\mathbf{z}_{n_{1} \to n_{2}} | \boldsymbol{\phi}_{n_{1} \to n_{2}}) q_{\text{mul}}(\mathbf{z}_{n_{1} \leftarrow n_{2}} | \boldsymbol{\phi}_{n_{1} \leftarrow n_{2}}).$$
(1)

where  $\{\gamma\}$ ,  $\{\lambda\}$ , and  $\{\phi\}$  are the variational parameters. Similarly, for the distributions over the model's parameters, we also use factorized  $q(\cdot)$  distributions, which are  $q_{\text{beta}}(\beta_k | \tau_{k1}, \tau_{k0})$  and  $q_{\text{dir}}(\boldsymbol{\omega}_i | \boldsymbol{\rho}_i)$ .

Given the factorized variational distribution  $q(\cdot)$ , we next bound the log likelihood of the observed data using Jensen's inequality. Specifically, we consider the so called evidence lower bound (ELBO) defined as follows:

$$\log p(\mathbf{Y}, \mathbf{w}_{1:N} | \boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) \geq \mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\gamma}, \boldsymbol{\lambda})$$
  
$$\triangleq \mathbb{E}_q[\log p(\mathbf{Y}, \mathbf{w}_{1:N}, \boldsymbol{\theta}_{1:N}, \mathbb{Z}_{1:N}, \mathbf{B}, \mathbf{C}, \Omega | \boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\kappa})] - \mathbb{E}_q[\log q(\boldsymbol{\theta}_{1:N}, \mathbb{Z}_{1:N}, \mathbf{B}, \mathbf{C}, \Omega)], \quad (2)$$

where we defined  $\Omega$  as a collection of  $\boldsymbol{\omega}.$ 

The ELBO in Equation 2 is expanded as follows:

$$\mathcal{L} = \sum_{n_1, n_2} \mathbb{E}_q[\log p(y(n_1, n_2) | \mathbf{z}_{n_1 \to n_2}, \mathbf{z}_{n_1 \leftarrow n_2}, \mathsf{B}] \\ + \sum_{n_1, n_2} \mathbb{E}_q[\log p(\mathbf{z}_{n_1 \to n_2} | \boldsymbol{\theta}_{n_1}) + \log p(\mathbf{z}_{n_1 \leftarrow n_2} | \boldsymbol{\theta}_{n_2})] \\ + \sum_n \sum_{m=1:M_n} \mathbb{E}_q[\log p(\mathbf{w}_{n,m}) | \mathbf{c}_m^n), \Omega] \\ + \sum_n \sum_{m=1:M_n} \mathbb{E}_q[\log p(\mathbf{c}_m^n | \mathbf{z}_{n \to \cdots}, \mathbf{z}_{\cdot \leftarrow n})] \\ + \sum_n \mathbb{E}_q[\log p(\boldsymbol{\theta}_n | \boldsymbol{\alpha})] \\ + \sum_k \mathbb{E}_q[\log p(\beta_k | \boldsymbol{\eta})] + \sum_k \mathbb{E}_q[\log p(\boldsymbol{\omega}_k | \boldsymbol{\kappa})] \\ - \sum_{n_1, n_2} \mathbb{E}_q[\log q(\mathbf{z}_{n_1 \to n_2} | \boldsymbol{\phi}_{n_1 \to n_2}) + \log q(\mathbf{z}_{n_1 \leftarrow n_2} | \boldsymbol{\phi}_{n_1 \leftarrow n_2})] \\ - \sum_n \sum_{m=1:M_n} \mathbb{E}_q[\log q(\mathbf{c}_m^n | \boldsymbol{\lambda}_m^n)] \\ - \sum_n \mathbb{E}_q[\log q(\boldsymbol{\theta}_n | \boldsymbol{\gamma}_n)].$$
(3)

The lower bound can be maximized using the coordinate ascent algorithm. Toward that goal, we take the (partial) derivatives of  $\mathcal{L}(\phi, \gamma, \lambda)$  with respect to the variational parameters and set them to zero. As a result, we obtain a set of iterative update equations for the parameters, whose fixed points correspond to local optima of the objective function.

For  $\boldsymbol{\gamma}_n$ , the update equations are as follows:

$$\gamma_{n,k} \leftarrow \alpha_k + \sum_{n' \neq n} \phi_{n \to n',k} + \sum_{n' \neq n} \phi_{n' \leftarrow n,k}.$$
 (4)

Similar reasoning leads to update equations for the parameters  $\{\phi\}$  and  $\{\lambda\}$ , with additional constraints that each component of those vectors sum to 1. As we mentioned above, here we are interested in undirected social interactions, which reduces the number of parameters due to the equalities  $\phi_{n \leftarrow n'} = \phi_{n' \rightarrow n}$ . Thus, we consider only the update equations for  $\phi_{n \rightarrow n'}$ .

After adding the appropriate Lagrange multipliers to ensure the normalization of the components,  $\sum_{n'} \phi_{n \to n'} = 1$ , we isolate the terms that contain  $\phi_{n \to n'}$ , and set their derivatives with respect to  $\phi_{n \to n'}$  to zero. The corresponding update equations are as follows:

$$\phi_{n \to n',k} \propto \exp(\mathbb{E}_q[\log p(\boldsymbol{\theta}_{n,k})] + \mathbb{E}_q[\log p(y(n,n'))] + \mathbb{E}_q[\log \prod_{m=1}^{M_n} p(\mathbf{w}_{n,m})^{\mathbf{c}_{m,n'}^n}]), \quad (5)$$

$$\boldsymbol{\lambda}_{m,n'}^n \propto \exp(\mathbb{E}_q[\log p(\mathbf{w}_{n,m}|\mathbf{c}_{m,n'}^n=1)]), \qquad (6)$$

where the term n' appearing in  $\mathbf{c}_{m,n'}^n$  and  $\boldsymbol{\lambda}_{m,n'}^n$  denotes the index of  $\phi_{n\to n'}$  in  $\mathbb{Z}_n$ . As with the indicator vector  $\mathbf{z}_{n\to n'}$ ,  $\mathbf{c}_m^n$  should have only one component equal to 1 setting all others to 0.

Note that the number of the parameters  $\phi_{n \to n'}$ , and thus the computational complexity of the update equations for Eq. 5, is quadratic in the number of nodes, even when the network is sparse. This is because the parameters are defined both for links and non-links. To avoid this computational bottleneck, we adopt an approximation technique [28], by assuming that the parameters  $\phi_{n\to}$ . for non-links can be replaced by a single mean-field parameter. Namely, let  $\{\phi_{n\to},\}^+$  and  $\{\phi_{n\to},\}^-$  be the set of parameters for links and non-links, respectively, and let  $\bar{\phi}_{n\to}$ . be the average over the set  $\{\phi_{n\to},\}^+$ . Within the above mean field approximation, each element of  $\{\phi_{n\to},\}^-$  is replaced by  $\bar{\phi}_{n\to}$ . Thus, the time complexity of algorithm becomes linear in the number of existing links in the network.

Further gains in computational efficiency is achieved by limiting the number of components in set  $\mathbb{Z}_n$  to the number of edges incident on node n, rather than having all the relations in the set. This corresponds to reusing indicator variables only when they have generated a link. Note also that now Equation 4 does not require all the parameters  $\{\phi\}$ , as the parameters corresponding to non-links are replaced by  $\bar{\phi}_{n\to \cdots}$ .

Finally, the update Equation 5 can be further simplified by making the best use of the assortative property in aMMSB, where only diagonal components of the block matrix are being considered. Instead of updating  $K \times K$  combinations of  $\phi_{n_1 \rightarrow n_2}$  and  $\phi_{n_2 \rightarrow n_1}$  for a link between node  $n_1$  and  $n_2$ , we can update only K parameters of by disregarding the inter-community links. For our aMMSB model, we use the following update equation:

$$\boldsymbol{\phi}_{n_1 \to n_2, k} \propto \exp(\mathbb{E}_q[\log p(\boldsymbol{\theta}_{n_1, k})] + \mathbb{E}_q[\log p(\boldsymbol{\theta}_{n_2, k})] + \mathbb{E}_q[\beta_k] + \mathbb{E}_q[\log \prod_{m=1}^{M_{n_1}} p(\mathbf{w}_{n_1, m})^{\mathbf{c}_{m, n_2}^{n_1}}]), \quad (7)$$

where we further use  $\mathbb{E}_q[\log p(\boldsymbol{\theta}_{n,k})] = \psi(\boldsymbol{\gamma}_{n,k}) - \psi(\sum_r \boldsymbol{\gamma}_{n,r}),$  $\mathbb{E}_q[\boldsymbol{\beta}_k] = \psi(\boldsymbol{\eta}_{k,1}) - \psi(\boldsymbol{\eta}_{k,2})$  employing the exponential family distribution property. As for the last term in Equation 7, we use the following equation:

$$\mathbb{E}_{q}[\log \prod_{m=1}^{M_{n_{1}}} p(\mathbf{w}_{n_{1},m})^{\mathbf{c}_{m,n_{2}}^{n_{1}}}] = \sum_{m=1}^{M_{n_{1}}} \lambda_{m,n_{2}}(\psi(\sum_{r} \mathbf{1}(\mathbf{w}_{n_{1},m}=i)\rho_{k,r}) - \psi(\sum_{i} \rho_{k,i})).$$
(8)

Let us now focus on variational distributions over the model parameters  $\beta_k$  and  $\omega_k$ , k = 1, ..., K. We had previously defined a Beta distribution  $q_{\text{beta}}(\beta_k | \tau_{k1}, \tau_{k0})$  with variational parameter  $\tau_1, \tau_0$ , and a Dirichlet distribution  $q_{\text{dir}}(\omega_i | \boldsymbol{\rho}_i)$  with variational parameter  $\boldsymbol{\rho}_k$ . Here we omit the derivation details and only present the final update equations for these parameters:

$$\boldsymbol{\tau}_{k,1} \leftarrow \boldsymbol{\eta}_1 + \sum_{(n_1,n_2) \in \text{link}} \boldsymbol{\phi}_{n_1 \to n_2,k},$$
 (9)

$$\boldsymbol{\tau}_{k,2} \leftarrow \boldsymbol{\eta}_{2} + \sum_{(n_{1},n_{2})\in\text{non-link}} \boldsymbol{\phi}_{n_{1}\rightarrow n_{2},k} \boldsymbol{\phi}_{n_{2}\rightarrow n_{1},k},$$

$$\boldsymbol{\rho}_{ij} \leftarrow \boldsymbol{\kappa}_{j} + \sum_{n=1}^{N} \sum_{m=1}^{M_{n}} \mathbf{1}(\mathbf{w}_{n,m} = j) \sum_{n'\in\text{link}} \boldsymbol{\phi}_{nn',i} \boldsymbol{\lambda}_{m,n'}.$$

$$(10)$$

Once the variational parameters are found, we can use them to estimate the model parameters themselves. We note that as an alternative approach, one can also derive explicit update equations for the parameter  $\omega_i$  directly, without using the variational parameter  $\rho_i$  in Equation 10.  $\omega_i$  can be optimized by introducing a Lagrange multiplier where we have the update equation as follows:

$$\boldsymbol{\omega}_{ij} \propto \sum_{n=1}^{N} \sum_{m=1}^{M_n} \mathbf{1}(\mathbf{w}_{n,m} = j) \sum_{n' \in \text{link}} \boldsymbol{\phi}_{nn',i} \boldsymbol{\lambda}_{m,n'}.$$
(11)

However, for this case, extra caution is needed that guarantees non-zero entities in any  $\omega_i$ . This can be easily achieved by using smoothing techniques such as Laplace smoothing or pseudo-count smoothing [35]. Our experimental results reflect the latter smoothing technique.

# 4. EXPERIMENTS

In this section we perform an extensive experimental evaluation of the Constrained Latent Space Model (CLSM). We first discuss some experiments performed with synthetic data, designed to better understand the main differences between our method and the Relational Topic Model (RTM). Then, we carry out an extensive experimental analysis using a number of real-world datasets from various technosocial systems. After describing the characteristics of these datasets, we present two problems, link and behavior prediction, and demonstrate the superiority of the proposed model over existing methods. We conclude discussing two case studies and the unique insights yielded by CLSM's use.

#### 4.1 Experiments with Synthetic Data

The main advantage of CLSM over RTM is that our model uses the aMMSB whereas RTM employs a simplified mechanism for the link generation process. We thus expect that CLSM will prove effective in presence of nodes with wellmixed memberships (or topics). By incorporating aMMSB, our model can better infer mixed topics while relying less on attribute information, if compared to RTM which heavily relies on users' attributes for accurate inference. In fact, in MMSB each interaction is captured to describe the topic distribution of a given node, while RTM simply posits a single distribution and tries to describe every interaction of a given node with that single distribution.

Figure 2 summarizes the experiments that confirm this intuition. Let us consider nodes with two topics where some of the nodes are exclusively assigned to one topic (or community) and others are well mixed between the two.<sup>3</sup> In real social systems, two distinct communities can exhibit similar or very different behaviors. To reflect this, we perform two experiments: in the former the two communities are well separated (Figure 2a, left panel), while in the latter we impose some arbitrarily-tuned degree of similarity between them (Figure 2a, right panel). The community-to-behavior distribution (equivalent to the per-topic word distribution) for the two cases has been controlled as in Figure 2a.

To measure the performance of CLSM and RTM, in Figure 2b we plot the *mean absolute error* (MAE) [31] of reconstruction against the gap between the two distribution peaks (the larger, the less similar are the two distributions). In the experiments, we generated synthetic data for 2,500 nodes and 1,500 behaviors (words) with two communities (topics) using the generative process of RTM. Figure 2b shows that, as the gap between the two per-topic word distributions tend to close, the performance of RTM starts to degrade, whereas the performance of our model remains substantially unchanged.

Further inspection revealed that, when the two topic distributions become similar one other, RTM tends to assign every node to the two extreme topics (i.e., pure topic1 or pure topic2), even when the nodes actually exhibit a clear mixture of the two topics. This causes the RTM's higher MAE when the two word distributions get similar, whereas CLSM consistently infers topics in line with the ground truth for all nodes including those exhibiting well-mixed topics.

#### 4.2 Experiments with Real-World Data

In the following we present experiments with various realworld datasets. Our aim is twofold: first and foremost, we want to illustrate the flexibility of our framework to adapt and perform well within various learning tasks commonly occurring when analyzing multi-modal techno-social systems; second, we want to prove that CLSM outperforms the state of the art. We will use data from diverse types of social systems, spanning various social networks, review sites, purchasing platforms, etc. The datasets are presented in §4.2.2.

#### 4.2.1 Setting the Performance Evaluation

In our experiments, we evaluate our model against RTM and Pairwise Link-LDA on two tasks: link prediction ( $\S4.2.3$ ), and attributes prediction ( $\S4.2.4$ ). We set Pairwise Link-LDA to use aMMSB (rather than fMMSB as in the original paper [45]) as this yields better accuracy, especially in the link prediction task. Through this fair comparison, we show that the improved performance of our method is not only the result of using aMMSB for mixed memberships, which consistently outperforms the original fMMSB.

For both prediction tasks, we follow the settings in [12] and employ a 5-fold cross validation approach: we use the



Figure 2: (a) Representation of two topics as distribution over the words: the left and right panels indicate weaklyoverlapping and strongly-overlapping topics, respectively. The difference between the topics can be measured by the distance between their respective peaks. (b) Mean absolute error of inference as a function of topic overlap (measured by the distance between the peaks).

4/5 of the data to train our models and then test their performance using the other 1/5. For link prediction, we hide all the links attached to the nodes in the test set, and compute the probability of a link from all the nodes in the test set to all the nodes in the training set. Similarly, for attribute prediction, we hide all the attributes of nodes in the test set, and compute the probability of each attribute. For the experiments below, we set the precision of the Dirichlet prior  $\alpha$  to 1,  $\epsilon$  as low as 1e-30, and run the Variational EM algorithm until its convergence. For the convergence criterion, the algorithm stops when the proportional change in the likelihood bound is less than 1e-8.

#### 4.2.2 Description of the Datasets

We here present the various datasets employed for performance evaluation. Some descriptive statistics are summarized in Table 1. All datasets provide a social network dimension. As for the attributes, *Gowalla* contains value counts, since the system allows for repeated behaviors (*checkins*) on the same attribute (venue). All the other datasets have binary representations over the attributes, which capture their presence (0 or 1). It's worth noting that, although the sheer size of these datasets is not massive due to the challenge of finding multi-modal social datasets for which both social network and user behaviors are available for large sets of users, we stress again how the low computational cost of CLSM makes it very suitable for large scale analysis.

Location-Based Social Network. *Gowalla* is a locationbased social network which allows users to *check-in* their current location using their mobile devices and share that with their friends. Cho *et al.* [14] collected the Gowalla check-in and social network data from Feb. 2009 to Oct. 2010. Each

<sup>&</sup>lt;sup>3</sup>We set Dirichlet prior  $\boldsymbol{\alpha} = [.3 \ .3]^{\top}$ 

Table 1: Summary statistics of the six multi-modal social datasets adopted for analysis and performance evaluation.

Dataset	# of Nodes	Size of Attributes	# of Behaviors	# of Links
Gowalla (San Francisco)	931 users	1,909 venues	63,543 check-ins	3,429 friendships
Ciao	1,442 users	2,238 products	49,210 word counts 46,732 reviews	29,040 user-trusts
Amazon Orkut Instagram	2,942 products 3,201 users	2,573 groups	29,009 descriptions 26,658 affiliations 01 246 tags	47,871 friendships

check-in datapoint consists of user ID, venue ID, timestamp, and location (latitude/longitude coordinates). In our experiments, we focused on the most represented US city: San Francisco. We consider users as nodes, friendships as links, and the number of check-ins on a given venue as attributes. Note that we discard any temporal information about the check-ins, and we focus on the top active 20% users, which yield over 80% of the total check-in records.

**Citation Network.** The *Cora* dataset [40, 51] was the largest citation network used to benchmark RTM's performance [12]. It contains the abstracts from the Cora research paper search engine, where the documents in this repository cite each other. We consider documents as nodes, citations as links, and the set of lexical terms as set of attributes. We adopt the pre-processed dataset used by Chang and Blei [12] where the observed attributes for each node convey the presence (0 or 1) of a given term in the paper abstract.

**Review-Trust Network.** This dataset consists of reviews on various products spanning DVDs to cars, and the trust network among *Ciao*'s users. Each review consists of a user (or reviewer), the product, the category of the product being reviewed, the review score, and the score that measures the helpfulness. To simplify, we disregard review and helpfulness scores, and assume that the presence of a review (regardless of the score) indicates a certain amount of interest in the product by the reviewer (e.g., which led the reviewer to buy the product and review it). Each user establishes a directed trust link with other users whom they want to follow. We treat those trust relationships as links, and the list of reviewed products as attributes.

**Co-purchase Network.** In this dataset, an undirected edge describes the "customers who bought this item also bought" relationship. We consider each product as node in our model. Each product also has a label description as a category, provided by *Amazon*. Most items are matched to multiple hierarchical label descriptions that correspond to the product. For instances, many digital cameras have multiple labels with Electronics> Camera & Photo> Digital Cameras, and are often purchased together with memory cards, tripods, or batteries. Here, we use these label descriptions as attributes.

Social Network with Special Interest Groups. Orkut is a social network service owned by Google. The dataset has been originally collected by Mislove *et al.* [42] during a crawl performed during Oct-Nov 2006. Each user in Orkut makes online friends with others and can also join *special interest groups*. Users are allowed to join as many groups as they would like. A group might consist of colleagues, celebrity fans, etc. Users can make new friends in a group, thus two users being in a same group does not necessarily imply that they are friends each other. Here, we consider the top 5,000 communities with the highest quality as identified by Yang and Leskovec [55]. Users with no connections to these communities have been filtered out. The dataset has been further refined by excluding the users with no links, and removing the groups with no nodes of our interest. In our experiments, we use the group affiliations as attributes and the friendships between users as links.

Media-Sharing Social Network. Instagram is a social multi-media sharing platform owned by Facebook. This dataset has been collected during Jan-Feb 2014, starting from 2,100 randomly-selected user seeds who participated to at least one of 72 popular photography contests, as identified by particular hashtags [22]. The social network mode reports follower relationships, whereas the behaviors represent the tags adopted by each user to label photos (attributes) chosen to participate to a given contest.

#### 4.2.3 Task 1: Link Prediction

Here we discuss our first evaluation task, link prediction. Accurate link prediction can be very useful in many realworld applications: for instance, in a location-based social network, we can use such tool to recommend friends to new users who have just joined the service and have not yet made any friends but already left some records of their activity.

Borrowing the language of topic modeling, using the four folds of data for training, (i) we first infer the topic distribution using the attributes, (ii) then compute the probability of the links between the nodes, and (iii) finally rank the pairs according to their likelihoods. These sorted candidates are compared with the ground truth, represented by the fifth fold left out for testing purposes. Our performance evaluation measure is the average ranking score. Figure 3 shows the average ranking of positive edges with 100% recalls. Lower average ranking scores imply better performance: a user would find their "real" friends higher up in the recommendation list. For each dataset, we increase the number of topics from 5 to 25 by increment of 5, and for each configuration we perform 10 rounds of cross validation and average the results. We also compare the performance of the benchmarked models using the Area Under the ROC Curve (AUC): the results are reported in Table 2.

For illustration purpose, let us discuss *Gowalla* first. Our approach performs better than the baselines (Pairwise Link-LDA and RTM) for any number of topics. As the number of topics increases, Pairwise Link-LDA fails to capture the joint topic space, whereas CLSM shows improvements in performance. Interestingly, when the number of topic is small, Pairwise Link-LDA performs better than RTM. We believe that a small number of topics functions as a constraint that enforces the two topic spaces to be close each other. CLSM



Figure 3: Average ranking score of link prediction on six multi-modal datasets. Lower scores  $\Leftrightarrow$  better performance.

clearly outperforms the baselines in terms of AUC scores, yielding a prediction accuracy of about 70%, a relative improvement of 15-18% over RTM and Pairwise Link-LDA.

We repeated the same experiments on the other datasets as well. Our model performs better on the link prediction task on every dataset and for all choices of topic sizes (see Figure 3). The link prediction results for Pairwise Link-LDA in Cora are better than those reported in previous studies [45] as we take advantage of aMMSB. Overall, our model along with RTM achieves better results as the number of topic increases — except for the Ciao dataset which exhibits an atypical pattern (the prediction accuracy gets worse as the number of topic increases). We believe this is due to the unique characteristic of the network in this platform, namely that a link reflects whether or not the user trusts another user as a reviewer. Many factors besides the topic we inferred may play a role in trust, such as the reputation of a reviewer: it is natural to trust a reviewer who submits compelling reviews although not necessarily sharing one's same interests. Otherwise, CLSM is the only model to consistently improve in performance as the number of topics increases.

In terms of AUC scores, CLSM yields performance around 75% for *Instagram*, *Cora* and *Ciao*, and better than 90% with *Amazon* and *Orkut*, exhibiting consistent and significant improvements. *Orkut* data benefit the least from CLSM with a modest 2.2% improvement, whereas *Instagram* data boast a remarkable 22% relative increment. For link prediction, CLSM yields an average increase across the six datasets of 12.56% and 9.06% respectively over Pairwise Link-LDA and RTM.



Figure 4: Average ranking score of attribute prediction on the six datasets. Lower scores  $\Leftrightarrow$  better performance.

# 4.2.4 Task 2: Attribute Prediction

In this second task, we predict user attributes based on the attributes of other users while leveraging network information. For each node in the test set, we hide its attributes but assume that its social network links are known. Similarly to link recommendation, such predictions can be used to suggest products (or more general attributes) to a user. As in the link prediction problem, the average ranking of true attributes with 100% recall is measured. The results are presented in Figure 4. Lower ranking scores imply that the actual attributes of a given node would appear ranked higher in the recommendation system.

Our model provides better predictions on all datasets, outperforming the two baselines on the attribute prediction task as well. The prediction accuracy scores (measured by AUC) for the three models are reported in Table 2: CLSM yields a performance improvement on average of 5.14% and 10.16% over Pairwise Link-LDA and RTM respectively. Our model is also the only one that exhibits a constant improvement as the number of topics grows: the average ranking scores produced by CLSM monotonically decreases in all datasets.

All models do well with *Amazon* data: a plausible hypothesis is that *Amazon* uses co-purchasing patterns to perform recommendations, and therefore the high AUC scores capture the ability of all models to reproduce such feature. We also note that Pairwise Link-LDA performs better than RTM on *Gowalla* venues prediction, and on *Amazon* category prediction. This may be due to the flexibility that aMMSB allows in the presence of social network interactions where people tend to be assigned to multiple communities.

AUC scores with $K = 15$					
Dataset	Model	Link AUC	Attribute AUC		
	PL-LDA	0.6278	0.7154		
Gowalla (SF)	RTM	0.6300	0.6746		
	CLSM	0.7010	0.7429		
	PL-LDA	0.6341	0.7928		
Cora	RTM	0.6509	0.7971		
	CLSM	0.7570	0.8203		
	PL-LDA	0.7038	0.7636		
Ciao	RTM	0.6738	0.6850		
	CLSM	0.7532	0.7929		
	PL-LDA	0.8565	0.9661		
Amazon	RTM	0.9586	0.9752		
	CLSM	0.9686	0.9765		
	PL-LDA	0.8837	0.6884		
Orkut	RTM	0.8836	0.6580		
	CLSM	0.9031	0.7890		
	PL-LDA	0.6385	0.7113		
Instagram	RTM	0.6971	0.6598		
	CLSM	0.7791	0.7401		
Instagram	RTM CLSM	0.6971 0.7791	0.6598 <b>0.7401</b>		

Table 2: Accuracy (AUC) in link and attribute prediction.

The prediction tasks we presented also unveiled some limits and challenges of these models: for example, the improvement aMMSB yields tend to reduce in *Cora*, possibly as a consequence of the presence of few topics in the documents. Finally, *Gowalla* data suggest that location-based social networks pose some serious challenges, yielding an accuracy of only 70% and 75% for link and attribute prediction respectively. We will explore in further details these two scenarios in the following in-depth analysis.

# 4.3 Case Studies

Case Study 1: Analysis of Cora. In this data set, each document is assigned to one of the following categories: Neural Networks, Rule Learning, Reinforcement Learning, Probabilistic Methods, Theory, Genetic Algorithms, and Case Based in Machine Learning. The assignment was performed manually by the dataset's curators [40, 51]. Figure 5a shows the total number of documents for each category: over half of the documents belong to either Neural Networks, Probabilistic Methods, or Genetic Algorithms. Without access to this category information, our model only uses the contained words and the citation link information to infer the topics of the documents. We here investigate how these categories match the topics we inferred. For each document, we specify a major topic and collect all documents in the corpus related to that topic. Figure 5b shows the percentage of categories for each of the 10 topics, where we have used K = 10 (number of topics) for illustrative purposes. The documents which have topic 1 as major topic are dominated by Neural Networks. We can also see how many topics favor specific categories over others. Some of the topics contain mixtures of labels: this is due to the loose separation between the category labels. When the categories are closely related, this tends to be captured in the pie chart of our inferred topics forming a mixture. With higher number of total topics (e.g., K = 25), we can observe a better separation of labels over topics. However, even with small size of K, we can still obtain meaningful insights such as how different categories relate each other (i.e., whether two categories are often captured together in a topic).



Figure 5: (a) The percentage of known categories in *Cora*. (b) The percentage of the categories for the collection of documents of each topics (K = 10).



Figure 6: The top 30 attended venues from the four most popular clusters are displayed in different colors. The emergence of neighborhoods appear evident.

Case Study 2: Analysis of Gowalla. To the best of our knowledge, our work is one of the first attempts at finding the joint topic space in a location-based social network. Previous studies modeled social networks and user mobility patterns [14, 17, 50, 53]. For example, Wang *et al.* [53] measured the co-visitation frequency to infer social interactions. However, their model suffers from capturing ties between nodes which often exhibit no co-occurrence in locations within physical proximity. We previously introduced a latent space representation of venues to predict social ties [16] using *agglomerative information bottlenecks* to cluster venues, and compared it to other models including one based on LDA. However, that approach only clusters venues, while ideally one would prefer to cluster venues and



Figure 7: Maps of San Francisco's four most attended clusters, and their top 100 attended venues. The three leftmost maps show clusters that mostly capture physical proximity. However, the rightmost map shows that CLSM can also capture other dimensions: in such a case for example we can see many tourist spots emerging due to the social nature of human traveling.

users simultaneously. Here, CLSM applied to Gowalla data finds the joint latent space that simultaneously describes the friendship between users and their *check-in* behavior (i.e., venues they attend). Using the geo-coordinates provided for each venue, we further investigate the geographical distribution of popular venues. Even though we have not explicitly included the geo-coordinates in our inference algorithm, we may reasonably assume that geo-clusters (neighborhoods) should be reflected in our inferred topics (clusters), since some people make friends with others who live or work closeby. Our analysis focuses on San Francisco: in Figure 6 we produce the results of the attribute prediction task with CLSM by setting the total number of topics (clusters) to K = 25. After selecting the four most popular clusters, we visualize the overall top 30 most attended venues: we can appreciate how the four clusters appear geographically well separated. Let us analyze them more into detail: Figure 7 this time shows the top 100 popular venues for each of the four clusters above. From left to right, the first three maps show that indeed the visited venues are affected by physical distances, and we can observe the emergence of neighborhoods. However, the rightmost map shows that our model is also able to capture clusters not necessarily representing physical closeness: further inspection reveals that this cluster represents tourist spots that are often attended by San Francisco's visitors. No other model was able to capture simultaneously such diverse patterns: CSLM overcomes this limit by leveraging the multi-modal nature of social data.

# 5. CONCLUSION

Techno-social systems exhibit a remarkable amount of complexity, capturing not only the interactions among users but also their attributes and behaviors along multiple dimensions. For example, in location-based social networks users can be concurrently described by their mobility patterns, their activities, their preferences, and of course by their social links. In online social networks users produce and consume contents, link each other, join groups, etc. Recent studies illustrated that these modalities taken independently cannot capture the multiple facets of user activity and behavior. As a result, our ability to effectively model, design, analyze and improve such systems substantially depends on the possibility of leveraging the abundance of rich contextual information. In summary, very many practical problems commonly occurring when designing or analyzing socio-technical systems would greatly benefit from a multimodal modeling framework.

To address these challenges, in this paper we proposed the Constrained Latent Space Model (CLSM), which employs a multi-modal paradigm to simultaneously describe social network information and user behavioral data using a latent space representation. To describe the network generative process, CLSM leverages Mixed Membership Stochastic Blockmodels (MMSB) that captures mixed memberships, nodes that may belong to more than one community at the same time. The latent space is inferred via Latent Dirichlet Allocation (LDA). One remarkable characteristic of CLSM is that it introduces a constraint that enforces MMSB and LDA to overlap on the same latent space without loss of flexibility. To tame the algorithmic complexity of such a task, we suggested an efficient inference strategy based on Variational Expectation Maximization, which scales linearly with the size of the network. Experiments with synthetic data illustrate the advantage of such an approach.

To further show the flexibility of our framework, we designed two evaluation tasks inspired by prediction problems commonly occurring with real techno-social systems: (i) a link prediction experiment that aims at reconstructing the missing links among users given available attributes or behavior data; (ii) an attribute prediction experiment where we infer user attributes (or behaviors) leveraging social interaction data. A rigorous evaluation of our model against two existing approaches (Pairwise Link-LDA and Relational Topic Model) illustrates the superior performance of CLSM. The benchmark performed on six different multi-modal social datasets includes location-based social networks, social sharing platforms, etc. We reported the increments in prediction accuracy (measured by AUC) yielded by CLSM with respect to the state of the art on all the evaluation scenarios: CLSM outperforms Pairwise Link-LDA on average by 12.56% and 5.14% in link and attribute prediction, respectively, and it surpasses RTM by 9.06% and 10.16% on the same tasks. We further discuss two scenarios in details, describing the insights obtained adopting our multi-modal framework that would have not been otherwise possible.

The main appealing features of the proposed approach include its excellent scalability and its ability to handle various types of multi-modal datasets with relational information. It is worth noting that, although we here report only analyses of bimodal networks, CLSM can be adopted with data capturing any number of modalities. For such a reason, in the future we plan to use our model on datasets with even more modalities (e.g., multi-layered attributes, or combinations of different types of relational information between nodes).

#### Acknowledgments

This research was supported in part by DARPA grant No. W911NF–12–1–0034.

## 6. REFERENCES

- L. A. Adamic and N. Glance. The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop* on *Link discovery*, pages 36–43. ACM, 2005.
- [2] S. Agreste, P. De Meo, E. Ferrara, S. Piccolo, and A. Provetti. Analysis of a heterogeneous social network of humans and cultural objects. *IEEE Transactions on Systems, Man and Cybernetics: Systems*, 45(4):559–570, 2015.
- [3] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. J. Mach. Learn. Res., 9:1981–2014, June 2008.
- [4] F. Al Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. In Proc. 6th International AAAI Conference on Weblogs and Social Media (ICWSM), pages 387–390, 2012.
- [5] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528. ACM, 2012.
- [6] R. A. Baños, J. Borge-Holthoefer, and Y. Moreno. The role of hidden influentials in the diffusion of online information cascades. *EPJ Data Science*, 2(1):1–16, 2013.
- [7] D. M. Blei. Probabilistic topic models. Communications of the ACM, 55(4):77–84, 2012.
- [8] D. M. Blei and M. I. Jordan. Modeling annotated data. In Proceedings of the 26th Annual International ACM the IR Conference on Research and Development in Information Retrieval, 2003.
- [9] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. J. Mach. Learn. Res., 3:993–1022, Mar. 2003.
- [10] C. Budak, D. Agrawal, and A. El Abbadi. Limiting the spread of misinformation in social networks. In *Proceedings of the 20th international conference on World wide web*, pages 665–674. ACM, 2011.
- [11] M. Cha, F. Benevenuto, Y.-Y. Ahn, and K. P. Gummadi. Delayed information cascades in flickr: Measurement, analysis, and modeling. *Computer Networks*, 56(3):1066–1076, 2012.
- [12] J. Chang and D. M. Blei. Relational topic models for document networks. In Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, 2009.
- [13] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec. Can cascades be predicted? In Proceedings of the 23rd international conference on World wide web, pages 925–936. ACM, 2014.
- [14] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: User movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD* international conference on Knowledge discovery and data mining, 2011.
- [15] Y.-S. Cho, G. Ver Steeg, and A. Galstyan. Co-evolution of selection and influence in social

networks. In Proc. of the Twenty-Fifth Conference on Artificial Intelligence (AAAI-11), 2011.

- [16] Y.-S. Cho, G. Ver Steeg, and A. Galstyan. Socially relevant venue clustering from check-in data. In *KDD* Workshop on Mining and Learning with Graphs, 2013.
- [17] D. J. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg. Inferring social ties from geographic coincidences. *Proceedings of the National Academy of Sciences*, 107(52):22436–22441, Dec. 2010.
- [18] P. De Meo, E. Ferrara, F. Abel, L. Aroyo, and G.-J. Houben. Analyzing user behavior across social sharing environments. ACM Transactions on Intelligent Systems and Technology (TIST), 5(1):14, 2013.
- [19] P. De Meo, E. Ferrara, G. Fiumara, and A. Provetti. On facebook, most ties are weak. *Communications of the ACM*, 57(11):78–84, 2014.
- [20] P. A. Dow, L. A. Adamic, and A. Friggeri. The anatomy of large facebook cascades. In Proc. 7th International AAAI Conference on Weblogs and Social Media (ICWSM), pages 145–154, 2013.
- [21] E. Ferrara. A large-scale community structure analysis in facebook. *EPJ Data Science*, 1(9):1–30, 2012.
- [22] E. Ferrara, R. Interdonato, and A. Tagarelli. Online popularity and topical interests through the lens of instagram. In *Proceedings of the 25th ACM conference* on Hypertext and social media, pages 24–34. ACM, 2014.
- [23] E. Ferrara, O. Varol, F. Menczer, and A. Flammini. Traveling trends: social butterflies or frequent fliers? In Proceedings of the first ACM conference on Online social networks, pages 213–222. ACM, 2013.
- [24] E. Ferrara and Z. Yang. Quantifying the effect of sentiment on information diffusion in social media. *PeerJ Computer Science*, 1:e26, 2015.
- [25] A. Friggeri, L. A. Adamic, D. Eckles, and J. Cheng. Rumor cascades. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, pages 101–110, 2014.
- [26] S. Goel, D. J. Watts, and D. G. Goldstein. The structure of online diffusion networks. In *Proceedings* of the 13th ACM conference on electronic commerce, pages 623–638. ACM, 2012.
- [27] N. Z. Gong, A. Talwalkar, L. Mackey, L. Huang, E. C. R. Shin, E. Stefanov, E. R. Shi, and D. Song. Joint link prediction and attribute inference using a social-attribute network. ACM Transactions on Intelligent Systems and Technology (TIST), 5(2):27, 2014.
- [28] P. Gopalan, D. M. Mimno, S. Gerrish, M. J. Freedman, and D. M. Blei. Scalable inference of overlapping communities. In *Proceedings of the Advances in Neural Information Processing Systems* 25, 2012.
- [29] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal* of the American Statistical Association, 97:1090–1098, 2001.
- [30] B. Huberman, D. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1), 2008.

- [31] R. J. Hyndman and A. B. Koehler. Another look at measures of forecast accuracy. *International journal of forecasting*, 22(4):679–688, 2006.
- [32] J. Jiang, C. Wilson, X. Wang, W. Sha, P. Huang, Y. Dai, and B. Y. Zhao. Understanding latent interactions in online social networks. ACM Trans. Web, 7(4):18:1–18:39, 2013.
- [33] M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. An introduction to variational methods for graphical models. 1999.
- [34] K. Joseph, C. H. Tan, and K. M. Carley. Beyond "local", "categories" and "friends": clustering foursquare users with latent "topics". In Proceedings of the 2012 ACM Conference on Ubiquitous Computing, 2012.
- [35] D. Jurafsky and J. H. Martin. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1st edition, 2000.
- [36] M. Kim and J. Leskovec. Modeling social networks with node attributes using the multiplicative attribute graph model. In 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011), pages 400–409, 2011.
- [37] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013.
- [38] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In Proceedings of the 19th international conference on World wide web, pages 591–600. ACM, 2010.
- [39] K. Lerman and R. Ghosh. Information contagion: an empirical study of spread of news on digg and twitter social networks. In *Proceedings of 4th International Conference on Weblogs and Social Media (ICWSM)*, pages 90–97, 2010.
- [40] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3(2):127–163, 2000.
- [41] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [42] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM* SIGCOMM conference on Internet measurement, pages 29–42. ACM, 2007.
- [43] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *science*, 328(5980):876–878, 2010.
- [44] S. A. Myers and J. Leskovec. The bursty dynamics of the twitter information network. In *Proceedings of the*

23rd international conference on World wide web, pages 913–924. ACM, 2014.

- [45] R. M. Nallapati, A. Ahmed, E. P. Xing, and W. W. Cohen. Joint latent topic models for text and citations. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008.
- [46] M. Newman and A. Clauset. Structure and inference in annotated networks. arXiv preprint arXiv:1507.04001, 2015.
- [47] Y. Pang, Q. Hao, Y. Yuan, T. Hu, R. Cai, and L. Zhang. Summarizing tourist destinations by mining user-generated travelogues and photos. *Comput. Vis. Image Underst.*, 115(3):352–363, Mar. 2011.
- [48] T. P. Peixoto. Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. *Phys. Rev. E*, 92:042807, Oct 2015.
- [49] D. M. Romero, C. Tan, and J. Ugander. On the interplay between social and topical structure. Proc. 7th International AAAI Conference on Weblogs and Social Media (ICWSM), pages 516–525, 2013.
- [50] A. Sadilek, H. Kautz, and J. P. Bigham. Finding your friends and following them to where you are. In Proceedings of the 5th ACM International Conference on Web Search and Data Mining, 2012.
- [51] P. Sen, G. M. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
- [52] M. Szell, R. Lambiotte, and S. Thurner. Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Sciences*, 107(31):13636–13641, 2010.
- [53] D. Wang, D. Pedreschi, C. Song, F. Giannotti, and A.-L. Barabasi. Human mobility, social ties, and link prediction. In *Proceedings of the 17th ACM SIGKDD* international conference on Knowledge discovery and data mining, 2011.
- [54] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *Proceedings of* the 20th international conference on World wide web, pages 705–714. ACM, 2011.
- [55] J. Yang and J. Leskovec. Defining and evaluating network communities based on ground-truth. In *ICDM*, pages 745–754. IEEE Computer Society, 2012.
- [56] J. Yang, J. McAuley, and J. Leskovec. Community detection in networks with node attributes. In 2013 IEEE 13th International Conference on Data Mining (ICDM), pages 1151–1156, 2013.
- [57] Y. Zhu, X. Yan, L. Getoor, and C. Moore. Scalable text and link analysis with mixed-topic link models. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '13, pages 473–481, New York, NY, USA, 2013. ACM.