

## MULTIMODAL FACTOR ANALYSIS

*Yasin Yilmaz and Alfred O. Hero*

University of Michigan, Ann Arbor

### ABSTRACT

A multimodal system with Poisson, Gaussian, and multinomial observations is considered. A generative graphical model that combines multiple modalities through common factor loadings is proposed. In this model, latent factors are like summary objects that has latent factor scores in each modality, and the observed objects are represented in terms of such summary objects. This potentially brings about a significant dimensionality reduction. It also naturally enables a powerful means of clustering based on a diverse set of observations. An expectation-maximization (EM) algorithm to find the model parameters is provided. The algorithm is tested on a Twitter dataset which consists of the counts and geographical coordinates of hashtag occurrences, together with the bag of words for each hashtag. The resultant factors successfully localizes the hashtags in all dimensions: counts, coordinates, topics. The algorithm is also extended to accommodate von Mises-Fisher distribution, which is used to model the spherical coordinates.

**Index Terms**— multimodal data fusion, unsupervised learning, graphical models, Twitter

### 1. INTRODUCTION

In complex systems a variety of observation modes (e.g., sensor readings, images, text) may be available to the decision maker. For instance, the emerging technologies, such as cyber-physical systems, internet of things, autonomous driving, and smart grid, can provide such a rich observation space, both in volume and modality. In monitoring a complex system, all modalities bear information about the system's internal state. In some cases, an anomaly may not be detected in each modality alone, but can be easily detected through a joint processing.

Unsupervised learning methods are instrumental in discovering hidden structures in data (e.g., factor analysis [1], topic modeling [2]), which can then be used to perform dimensionality reduction, anomaly detection, and clustering. Conventionally, they deal with unimodal data [1, 2]. Efficient

fusion of multimodal data brings about information diversity and can greatly improve the statistical inference performance [3]. For example, data used for detection and estimation tasks are strongly coupled when both problems are solved jointly, considerably increasing the overall performance [4].

In [5], Gaussian and multinomial observations are jointly modeled using a mixture of factor analyzers. Although that work is conceptually similar to our work, its probabilistic model is quite different. In [5], the latent factor scores are shared in the system, which also synchronizes the modalities. Here we do not impose a synchronized observation model. We let the different modalities run in their own continuous-valued state, and only link them through the factor loadings of objects. On the contrary, in [5], modalities differ in their factor loadings, which can take a discrete set of values.

### 2. PROBLEM STATEMENT

We consider  $P$  objects (e.g., documents) each of which is observed through three disparate information sources. In this paper, we assume Poisson, Gaussian, and multinomial statistical models for the information sources because of their wide application areas. Specifically, Poisson distribution is used to model event counts; Gaussian distribution is used for continuous-valued observations; and multinomial distribution models categorical observations. (Real-world examples can be seen in the Experiments section.)

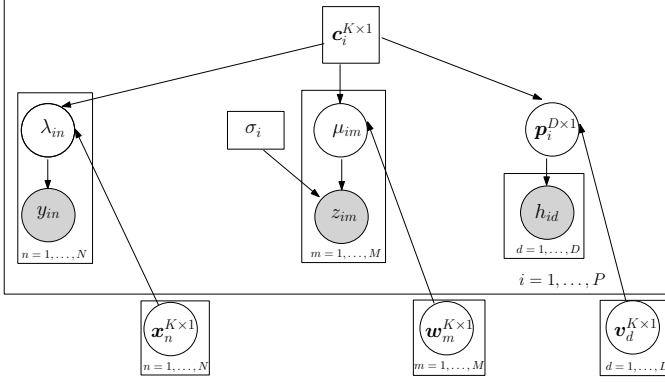
The graphical model in Fig. 1 is assumed to generate the multimodal observations  $\mathbf{Y}^{P \times N}$ ,  $\mathbf{Z}^{P \times M}$ , and  $\mathbf{H}^{P \times D}$ . For each object  $i$ , the latent factor scores  $\mathbf{x}_n$ ,  $\mathbf{w}_m$ , and  $\mathbf{v}_d$  for  $K$  factors are mixed in the the natural parameters of the distributions through the unknown factor loadings  $\mathbf{c}_i$ . Gaussian priors are assumed for the latent factor scores. In particular, each Poisson observation  $y_{in}$  is conditionally distributed as

$$\begin{aligned} y_{in} | \mathbf{x}_n &\sim \text{Pois}(e^{\mathbf{c}_i^T \mathbf{x}_n}), \\ \mathbf{x}_n &\sim \mathcal{N}(\boldsymbol{\zeta}, \mathbf{R}) \end{aligned} \quad (1)$$

for  $i = 1, \dots, P$ ,  $n = 1, \dots, N$ . Similarly, for each Gaussian observation  $z_{im}$  we have

$$\begin{aligned} z_{im} | \mathbf{w}_m &\sim \mathcal{N}(\mathbf{c}_i^T \mathbf{w}_m, \sigma_i^2), \\ \mathbf{w}_m &\sim \mathcal{N}(\boldsymbol{\alpha}, \mathbf{S}) \end{aligned} \quad (2)$$

This work was supported in part by the Consortium for Verification Technology under Department of Energy National Nuclear Security Administration award number DE-NA0002534, and the Army Research Office (ARO) grant number W911NF-11-1-0391.



**Fig. 1.** Generative graphical model. Plate representation is used to show repeated structures. Circles and rectangles represent random and deterministic variables, respectively. Observed variables are shaded. Each object  $i$  has three disparate observation streams: Poisson  $\{y_{in}\}_n$ , Gaussian  $\{z_{im}\}_m$ , and multinomial  $\{h_{id}\}_d$ . Factor loadings  $c_i$ , and the latent factor scores  $\mathbf{x}_n$ ,  $\mathbf{w}_m$ , and  $\mathbf{v}_d$  constitute the multimodal factor model.

for  $m = 1, \dots, M$ . We use different observation indices for different information sources since they do not need to be synchronized and the total number of observations may vary. For each multinomial observation  $h_{id}$  we have

$$h_{id}|\{\mathbf{v}_d\} \sim \text{Mult} \left( L_i; \frac{e^{c_i^T \mathbf{v}_1}}{\sum_{d=1}^D e^{c_i^T \mathbf{v}_d}}, \dots, \frac{e^{c_i^T \mathbf{v}_D}}{\sum_{d=1}^D e^{c_i^T \mathbf{v}_d}} \right), \quad (3)$$

$$\mathbf{v}_d \sim \mathcal{N}(\boldsymbol{\beta}, \mathbf{Q}),$$

where  $L_i$  is the total number of instances and  $h_{id}$  is the number of instances observed under category  $d$ .

The above probabilistic models are similar to generalized linear models since mixing occurs in the natural parameters. However, here the regressors  $\mathbf{x}_n$ ,  $\mathbf{w}_m$ , and  $\mathbf{v}_d$  are unknown, as well as the regression coefficients  $c_i$ , as opposed to generalized linear models.

### 3. EXPECTATION-MAXIMIZATION ALGORITHM

In this section, we derive the expectation-maximization (EM) algorithm to find the parameters

$$\theta = \{c_i, \boldsymbol{\zeta}, \mathbf{R}, \sigma_i^2, \boldsymbol{\alpha}, \mathbf{S}, \boldsymbol{\beta}, \mathbf{Q}\}.$$

#### 3.1. Poisson E-step

We are interested in computing the expectation of the complete-data log-likelihood  $\mathbb{E}[\log P(\{\mathbf{y}_n, \mathbf{x}_n\}|\theta)]$  over the posterior distribution  $P(\{\mathbf{x}_n\}|\{\mathbf{y}_n\}, \theta)$  of the latent factor scores. Due to the lack of conjugacy between the prior on  $\mathbf{x}_n$  and the Poisson likelihood, the posterior distribution does not have a closed

form expression. Therefore, as in [6], we approximate the posterior with a Gaussian whose mean and covariance are the mode of the posterior (i.e., MAP estimate of  $\mathbf{x}_n$ ) and the negative inverse Hessian of the log posterior at that mode.

We start by writing

$$\begin{aligned} \log P(\mathbf{x}_n|\mathbf{y}_n, \theta) &= \log P(\mathbf{y}_n|\mathbf{x}_n, \theta) + \log P(\mathbf{x}_n|\theta) + C_1 \\ &= \left( \sum_{i=1}^P -e^{c_i^T \mathbf{x}_n} + y_{in} c_i^T \mathbf{x}_n \right) \\ &\quad - \frac{1}{2} \mathbf{x}_n^T \mathbf{R}^{-1} \mathbf{x}_n + \boldsymbol{\zeta}^T \mathbf{R}^{-1} \mathbf{x}_n + C_2, \end{aligned}$$

where  $C_1$  and  $C_2$  are constants that do not depend on  $\mathbf{x}_n$ , and we used the Poisson likelihood from (1) and the multivariate Gaussian pdf. Then, the gradient and the Hessian are given by

$$\begin{aligned} \nabla_{\mathbf{x}_n} \log P(\mathbf{x}_n|\mathbf{y}_n, \theta) &= \sum_{i=1}^P \left( -e^{c_i^T \mathbf{x}_n} + y_{in} \right) c_i - \mathbf{R}^{-1} (\mathbf{x}_n - \boldsymbol{\zeta}) \\ \nabla_{\mathbf{x}_n}^2 \log P(\mathbf{x}_n|\mathbf{y}_n, \theta) &= \sum_{i=1}^P -e^{c_i^T \mathbf{x}_n} c_i c_i^T - \mathbf{R}^{-1}. \quad (4) \end{aligned}$$

Since  $\log P(\mathbf{x}_n|\mathbf{y}_n, \theta)$  is strictly concave in  $\mathbf{x}_n$ , we can find the unique mode  $\boldsymbol{\xi}_n$  using the gradient and the Hessian in Newton's method, i.e.,

$$\mathcal{H}_{\mathbf{x}_n^t} (\mathbf{x}_n^{t+1} - \mathbf{x}_n^t) = -\mathcal{G}_{\mathbf{x}_n^t}, \quad \mathbf{x}_n^t \rightarrow \boldsymbol{\xi}_n,$$

where  $\mathcal{G}_{\mathbf{x}_n^t}$  and  $\mathcal{H}_{\mathbf{x}_n^t}$  are the gradient and the Hessian at iteration  $t$ , computed using (4).

We approximate the posterior as  $P(\mathbf{x}_n|\mathbf{y}_n, \theta) \approx \mathcal{N}(\boldsymbol{\xi}_n, \boldsymbol{\Psi}_n)$ , where  $\boldsymbol{\Psi}_n = -\mathcal{H}_{\boldsymbol{\xi}_n}^{-1}$ . Finally, the expected complete-data log-likelihood is given by

$$\begin{aligned} \mathbb{E}[\log P(\{\mathbf{y}_n, \mathbf{x}_n\}|\theta)] &= \sum_{n=1}^N \mathbb{E}[\log P(\mathbf{y}_n|\mathbf{x}_n, \theta) + \log P(\mathbf{x}_n|\theta)] \\ &= \sum_{n=1}^N \sum_{i=1}^P \mathbb{E} \left[ -e^{c_i^T \mathbf{x}_n} + y_{in} c_i^T \mathbf{x}_n \right] - \frac{1}{2} \mathbb{E}[\mathbf{x}_n^T \mathbf{R}^{-1} \mathbf{x}_n] \\ &\quad + \boldsymbol{\zeta}^T \mathbf{R}^{-1} \mathbb{E}[\mathbf{x}_n] - \frac{1}{2} \boldsymbol{\zeta}^T \mathbf{R}^{-1} \boldsymbol{\zeta} - \frac{1}{2} |\mathbf{R}| + C_3 \\ &\approx \sum_{n=1}^N \sum_{i=1}^P \left( -\mathbb{E} \left[ e^{c_i^T \mathbf{x}_n} \right] + y_{in} c_i^T \boldsymbol{\xi}_n \right) \\ &\quad - \frac{1}{2} \text{Tr}(\mathbf{R}^{-1} (\boldsymbol{\Psi}_n + \boldsymbol{\xi}_n \boldsymbol{\xi}_n^T)) + \boldsymbol{\zeta}^T \mathbf{R}^{-1} \boldsymbol{\xi}_n \\ &\quad - \frac{1}{2} \boldsymbol{\zeta}^T \mathbf{R}^{-1} \boldsymbol{\zeta} - \frac{1}{2} |\mathbf{R}| + C_3, \quad (5) \end{aligned}$$

where  $C_3$  is a constant and we used the fact that

$$\begin{aligned} \mathbb{E} [\mathbf{x}_n^T \mathbf{R}^{-1} \mathbf{x}_n] &= \mathbb{E} [\text{Tr}(\mathbf{x}_n^T \mathbf{R}^{-1} \mathbf{x}_n)] = \mathbb{E} [\text{Tr}(\mathbf{R}^{-1} \mathbf{x}_n \mathbf{x}_n^T)] \\ &= \text{Tr}(\mathbf{R}^{-1} \mathbb{E} [\mathbf{x}_n \mathbf{x}_n^T]) \\ &\approx \text{Tr}(\mathbf{R}^{-1} (\boldsymbol{\Psi}_n + \boldsymbol{\xi}_n \boldsymbol{\xi}_n^T)). \end{aligned} \quad (6)$$

### 3.2. Poisson M-step

In iteration  $t + 1$ , we find the parameters  $\boldsymbol{\zeta}^{t+1}$  and  $\mathbf{R}^{t+1}$  that maximize (5). Particularly,

$$\boldsymbol{\zeta}^{t+1} = \arg \max_{\boldsymbol{\zeta}} \boldsymbol{\zeta}^T \mathbf{R}^{-1} \sum_{n=1}^N \boldsymbol{\xi}_n^t - \frac{N}{2} \boldsymbol{\zeta}^T \mathbf{R}^{-1} \boldsymbol{\zeta}.$$

Equating the derivative to zero we find that the mean of the factor scores is given by the average of the posterior means, i.e.,

$$\boldsymbol{\zeta}^{t+1} = \frac{1}{N} \sum_{n=1}^N \boldsymbol{\xi}_n^t. \quad (7)$$

Similarly,

$$\begin{aligned} \mathbf{R}^{t+1} &= \arg \max_{\mathbf{R}} \sum_{n=1}^N -\frac{1}{2} \text{Tr}(\mathbf{R}^{-1} (\boldsymbol{\Psi}_n^t + \boldsymbol{\xi}_n^t (\boldsymbol{\xi}_n^t)^T)) \\ &\quad + \boldsymbol{\zeta}^T \mathbf{R}^{-1} \sum_{n=1}^N \boldsymbol{\xi}_n^t - \frac{N}{2} \boldsymbol{\zeta}^T \mathbf{R}^{-1} \boldsymbol{\zeta} - \frac{N}{2} |\mathbf{R}|. \end{aligned}$$

Taking the derivative and equating to zero we get

$$\begin{aligned} \sum_{n=1}^N (\mathbf{R}^{t+1})^{-1} \left[ \frac{1}{2} (\boldsymbol{\Psi}_n^t + \boldsymbol{\xi}_n^t (\boldsymbol{\xi}_n^t)^T) - \boldsymbol{\zeta}^{t+1} \boldsymbol{\xi}_n^T \right. \\ \left. + \frac{1}{2} \boldsymbol{\zeta}^{t+1} (\boldsymbol{\zeta}^{t+1})^T \right] (\mathbf{R}^{t+1})^{-1} - \frac{N}{2} (\mathbf{R}^{t+1})^{-1} = 0 \\ \mathbf{R}^{t+1} = \frac{1}{N} \sum_{n=1}^N (\boldsymbol{\Psi}_n^t + \boldsymbol{\xi}_n^t (\boldsymbol{\xi}_n^t)^T) - \boldsymbol{\zeta}^{t+1} (\boldsymbol{\zeta}^{t+1})^T. \end{aligned} \quad (8)$$

### 3.3. Gaussian E-step

We want to compute  $\mathbb{E} [\log \text{P}(\{z_m, \mathbf{w}_m\} | \theta)]$  over the posterior distribution  $\text{P}(\{z_m\} | \{\mathbf{w}_m\}, \theta)$  of the latent factor scores. Since the Gaussian prior on the factor scores is conjugate to the Gaussian likelihood, the posterior is also Gaussian. To find its mean and covariance, from (2), we write

$$\begin{aligned} \text{P}(z_m, \mathbf{w}_m | \theta) &= \text{P}(z_m | \mathbf{w}_m, \theta) \text{P}(\mathbf{w}_m | \theta) \\ &= \frac{e^{-\frac{1}{2} [(z_m - \mathbf{C} \mathbf{w}_m)^T \boldsymbol{\Sigma}^{-1} (z_m - \mathbf{C} \mathbf{w}_m) + (\mathbf{w}_m - \boldsymbol{\alpha})^T \mathbf{S}^{-1} (\mathbf{w}_m - \boldsymbol{\alpha})]}}{(2\pi)^{\frac{P+K}{2}} |\boldsymbol{\Sigma}|^{\frac{1}{2}} |\mathbf{S}|^{\frac{1}{2}}}, \end{aligned}$$

where  $\mathbf{C} = [c_1 \dots c_P]^T$  and  $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_P^2)$ . Collecting the terms that depend on  $z_m$  together and completing the

square we finally obtain the posterior mean and covariance as

$$\mathbf{a}_m = \mathbf{B} (\mathbf{C}^T \boldsymbol{\Sigma}^{-1} z_m + \mathbf{S}^{-1} \boldsymbol{\alpha}), \quad \mathbf{B} = (\mathbf{C}^T \boldsymbol{\Sigma}^{-1} \mathbf{C} + \mathbf{S}^{-1})^{-1}, \quad (9)$$

respectively.

Then, the expected complete-data log-likelihood is written as

$$\begin{aligned} &\mathbb{E} [\log \text{P}(\{z_m, \mathbf{w}_m\} | \theta)] \\ &= \sum_{m=1}^M \mathbb{E} [\log \text{P}(z_m | \mathbf{w}_m, \theta) + \log \text{P}(\mathbf{w}_m | \theta)] \\ &= -\frac{1}{2} \left\{ \sum_{m=1}^M \sum_{i=1}^P \left( \frac{1}{\sigma_i^2} \mathbb{E} [(z_{im} - c_i^T \mathbf{w}_m)^2] + \log(2\pi \sigma_i^2) \right) \right. \\ &\quad \left. + \mathbb{E} [(\mathbf{w}_m - \boldsymbol{\alpha})^T \mathbf{S}^{-1} (\mathbf{w}_m - \boldsymbol{\alpha})] + \log((2\pi)^K |\mathbf{S}|) \right\}. \end{aligned} \quad (10)$$

### 3.4. Gaussian M-step

At each iteration  $t + 1$ , we find the parameters  $\boldsymbol{\alpha}^{t+1}$ ,  $\mathbf{S}^{t+1}$ , and  $\{(\sigma_i^2)^{t+1}\}_i$  that maximize (10). Specifically,

$$\begin{aligned} \boldsymbol{\alpha}^{t+1} &= \arg \max_{\boldsymbol{\alpha}} -\frac{1}{2} \sum_{m=1}^M \mathbb{E} [(\mathbf{w}_m - \boldsymbol{\alpha})^T \mathbf{S}^{-1} (\mathbf{w}_m - \boldsymbol{\alpha})] \\ &= \arg \max_{\boldsymbol{\alpha}} \boldsymbol{\alpha}^T \mathbf{S}^{-1} \sum_{m=1}^M \mathbf{a}_m^t - \frac{M}{2} \boldsymbol{\alpha}^T \mathbf{S}^{-1} \boldsymbol{\alpha}, \end{aligned}$$

where  $\mathbf{a}_m^t$ , given in (9), is the posterior mean  $\mathbb{E}[\mathbf{w}_m]$  at iteration  $t$ . Taking the derivative we get

$$\boldsymbol{\alpha}^{t+1} = \frac{1}{M} \sum_{m=1}^M \mathbf{a}_m^t. \quad (11)$$

Similarly,

$$\begin{aligned} \mathbf{S}^{t+1} &= \arg \max_{\mathbf{S}} \\ &\quad -\frac{1}{2} \sum_{m=1}^M \mathbb{E} [(\mathbf{w}_m - \boldsymbol{\alpha})^T \mathbf{S}^{-1} (\mathbf{w}_m - \boldsymbol{\alpha})] - \frac{M}{2} \log |\mathbf{S}| \\ &= \arg \max_{\mathbf{S}} -\frac{1}{2} \sum_{m=1}^M \text{Tr}(\mathbf{S}^{-1} (\mathbf{B}_m + \mathbf{a}_m \mathbf{a}_m^T)) \\ &\quad + \boldsymbol{\alpha}^T \mathbf{S}^{-1} \sum_{m=1}^M \mathbf{a}_m^t - \frac{M}{2} \boldsymbol{\alpha}^T \mathbf{S}^{-1} \boldsymbol{\alpha} - \frac{M}{2} \log |\mathbf{S}|, \end{aligned}$$

where we used (6) to write  $\mathbb{E} [\mathbf{w}_m^T \mathbf{S}^{-1} \mathbf{w}_m]$ . Similar to (8), taking the derivative we find

$$\mathbf{S}^{t+1} = \frac{1}{M} \sum_{m=1}^M (\mathbf{B}_m^t + \mathbf{a}_m^t (\mathbf{a}_m^t)^T) - \boldsymbol{\alpha}^{t+1} (\boldsymbol{\alpha}^{t+1})^T. \quad (12)$$

We next find

$$(\sigma_i^2)^{t+1} = \arg \max_{\sigma_i^2} -\frac{1}{2\sigma_i^2} \sum_{m=1}^M \mathbb{E} [(z_{im} - \mathbf{c}_i^T \mathbf{w}_m)^2] - \frac{M}{2\sigma_i^2} \log \sigma_i^2,$$

where, from (9),  $\mathbf{c}_i^T \mathbf{w}_m | z_m \sim \mathcal{N}(\mathbf{c}_i^T \mathbf{a}_m, \mathbf{c}_i^T \mathbf{B} \mathbf{c}_i)$ , hence  $\mathbb{E} [(z_{im} - \mathbf{c}_i^T \mathbf{w}_m)^2] = \mathbf{c}_i^T \mathbf{B} \mathbf{c}_i + (\mathbf{c}_i^T \mathbf{a}_m - z_{im})^2$ . Equating the derivative to zero and solving for  $\sigma_i^2$  we get

$$(\sigma_i^2)^{t+1} = \frac{1}{M} \sum_{m=1}^M (z_{im} - \mathbf{c}_i^T \mathbf{a}_m)^2 + \mathbf{c}_i^T \mathbf{B} \mathbf{c}_i. \quad (13)$$

### 3.5. Multinomial E-step

In the multinomial case, for identifiability, we use the last category as pivot and write the likelihood in terms of the alternative factor scores  $\mathbf{u}_d = \mathbf{v}_d - \mathbf{v}_D$ ,  $d = 1, \dots, D$ ,

$$\begin{aligned} \mathbb{P}(\mathbf{h}_i | \{\mathbf{u}_d\}, \theta) &= \prod_{d=1}^D \left( \frac{e^{\mathbf{c}_i^T \mathbf{u}_d}}{1 + \sum_{l=1}^{D-1} e^{\mathbf{c}_i^T \mathbf{u}_l}} \right)^{h_{id}} \\ &= \prod_{d=1}^D e^{[\eta_{id} - \text{lse}(\boldsymbol{\eta}_i)] h_{id}}, \end{aligned} \quad (14)$$

where  $\eta_{id} = \mathbf{c}_i^T \mathbf{u}_d$  and  $\boldsymbol{\eta}_i = [\eta_{i1} \dots \eta_{iD-1}]$ . The normalizing term in the probability expression, also called the log-sum-exp (lse) function prevents a closed form solution for the posterior. Finding a quadratic upper bound for it we can bound from below the likelihood, and in turn the expected complete-data log-likelihood, which we want to maximize.

Using the Taylor series of  $\text{lse}(\boldsymbol{\eta}_i)$  we can find such a quadratic bound [7] as follows,

$$\begin{aligned} \text{lse}(\boldsymbol{\eta}_i) &= \text{lse}(\boldsymbol{\gamma}_i) + (\boldsymbol{\eta}_i - \boldsymbol{\gamma}_i)^T \nabla \text{lse}(\boldsymbol{\gamma}_i) \\ &\quad + \frac{1}{2} (\boldsymbol{\eta}_i - \boldsymbol{\gamma}_i)^T \nabla^2 \text{lse}(\boldsymbol{\gamma}_i) (\boldsymbol{\eta}_i - \boldsymbol{\gamma}_i) \\ &\leq \frac{1}{2} \boldsymbol{\eta}_i^T \mathbf{A} \boldsymbol{\eta}_i - \mathbf{b}_{\boldsymbol{\gamma}_i}^T \boldsymbol{\eta}_i + c_{\boldsymbol{\gamma}_i}. \end{aligned} \quad (15)$$

To show the above inequality note that  $\nabla \text{lse}(\boldsymbol{\gamma}_i)$  is the probability vector  $\mathbf{p}_i(\boldsymbol{\gamma}_i)$ , and  $\nabla^2 \text{lse} = \mathbf{\Lambda}_{\mathbf{p}_i} - \mathbf{p}_i \mathbf{p}_i^T$  where  $\mathbf{\Lambda}_{\mathbf{p}_i} = \text{diag}(p_{i1}, \dots, p_{iD-1})$ . In [7], the latter is shown to be bounded as follows

$$\nabla^2 \text{lse} \leq \mathbf{A} = \frac{1}{2} \left( \mathbf{I}_{D-1} - \frac{\mathbf{1}_{D-1} \mathbf{1}_{D-1}^T}{D} \right),$$

where  $\mathbf{I}_{D-1}$  and  $\mathbf{1}_{D-1}$  are the identity matrix and the vector of ones of size  $D-1 \times D-1$  and  $D-1 \times 1$ , respectively. Substituting  $\mathbf{A}$  and  $\nabla \text{lse}(\boldsymbol{\gamma}_i)$  and organizing the terms gives us the inequality in (15), where

$$\mathbf{b}_{\boldsymbol{\gamma}_i} = \mathbf{A} \boldsymbol{\gamma}_i - \mathbf{p}_i(\boldsymbol{\gamma}_i), \quad c_{\boldsymbol{\gamma}_i} = \text{lse}(\boldsymbol{\gamma}_i) + \frac{1}{2} \boldsymbol{\gamma}_i^T \mathbf{A} \boldsymbol{\gamma}_i - \mathbf{p}_i(\boldsymbol{\gamma}_i).$$

From (14) and (15), we get

$$\begin{aligned} \log \mathbb{P}(\mathbf{h}_i | \{\mathbf{u}_d\}, \theta) &\geq \mathbf{h}_i^T \boldsymbol{\eta}_i - \left( \frac{1}{2} \boldsymbol{\eta}_i^T \mathbf{A} \boldsymbol{\eta}_i - \mathbf{b}_{\boldsymbol{\gamma}_i}^T \boldsymbol{\eta}_i + c_{\boldsymbol{\gamma}_i} \right) L_i \\ &= -\frac{1}{2} \left( \boldsymbol{\eta}_i - \mathbf{A}^{-1} \left( \frac{\mathbf{h}_i}{L_i} + \mathbf{b}_{\boldsymbol{\gamma}_i} \right) \right)^T \mathbf{A} \\ &\quad \left( \boldsymbol{\eta}_i - \mathbf{A}^{-1} \left( \frac{\mathbf{h}_i}{L_i} + \mathbf{b}_{\boldsymbol{\gamma}_i} \right) \right) \\ &\quad + \frac{1}{2} \left( \frac{\mathbf{h}_i}{L_i} + \mathbf{b}_{\boldsymbol{\gamma}_i} \right)^T \mathbf{A}^{-1} \left( \frac{\mathbf{h}_i}{L_i} + \mathbf{b}_{\boldsymbol{\gamma}_i} \right) - c_{\boldsymbol{\gamma}_i}. \end{aligned}$$

Exponentiating we obtain the following lower bound for the likelihood

$$\mathbb{P}(\mathbf{h}_i | \{\mathbf{u}_d\}, \theta) \geq \mathcal{N}(\tilde{\mathbf{h}}_i | \boldsymbol{\eta}_i, \mathbf{A}^{-1}) f_i(\boldsymbol{\gamma}_i),$$

where  $\tilde{\mathbf{h}}_i = \mathbf{A}^{-1} \left( \frac{\mathbf{h}_i}{L_i} + \mathbf{b}_{\boldsymbol{\gamma}_i} \right) = \mathbf{A}^{-1} \left( \frac{\mathbf{h}_i}{L_i} - \mathbf{p}_i(\boldsymbol{\gamma}_i) \right) + \boldsymbol{\gamma}_i$  is the Gaussian pseudo-observation

Since the factor scores  $\{\mathbf{u}_d\}$  are correlated given the observations, we seek the posterior of the combined vector  $\mathbf{u} = [u_{11} \dots u_{D-1K}]^T \sim \mathcal{N}(\mathbf{0}, \tilde{\mathbf{Q}})$  where  $\mathbf{0}$  is the zero vector and  $\tilde{\mathbf{Q}} = \mathbf{I}_{D-1} \otimes \mathbf{Q}$  is a block-diagonal matrix. Similarly, defining  $\tilde{\mathbf{C}}_i = \mathbf{I}_{D-1} \otimes \mathbf{c}_i$  we can write  $\boldsymbol{\eta}_i = \tilde{\mathbf{C}}_i^T \mathbf{u}$ . Then, for the complete-data likelihood we have

$$\begin{aligned} \mathbb{P}(\{\mathbf{h}_i\}, \mathbf{u} | \theta) &\geq \\ &\left( \prod_{i=1}^P \mathcal{N}(\tilde{\mathbf{h}}_i | \tilde{\mathbf{C}}_i^T \mathbf{u}, \mathbf{A}^{-1}) f_i(\boldsymbol{\gamma}_i) \right) \mathcal{N}(\mathbf{u} | \mathbf{0}, \tilde{\mathbf{Q}}). \end{aligned} \quad (16)$$

Organizing the terms and completing the square we write

$$\mathbb{P}(\{\mathbf{h}_i\}, \mathbf{u} | \theta) \geq \mathcal{N}(\mathbf{u} | \boldsymbol{\phi}, \boldsymbol{\Phi}) g_i(\{\mathbf{h}_i, \boldsymbol{\gamma}_i\}),$$

where

$$\boldsymbol{\phi} = \boldsymbol{\Phi} \sum_{i=1}^P \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{h}}_i, \quad \boldsymbol{\Phi} = \left( \sum_{i=1}^P \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{C}}_i^T + \tilde{\mathbf{Q}}^{-1} \right)^{-1} \quad (17)$$

are the posterior mean and covariance.

### 3.6. Multinomial M-step

We maximize the expected complete-data log-likelihood of the lower bound given in (16) using the posterior mean and

covariance, given in (17), as in [5].

$$\begin{aligned}
\mathbf{Q}^{t+1} &= \arg \max_{\mathbf{Q}} \mathbb{E} [\log \mathbb{P}(\{\mathbf{h}_i\}, \mathbf{u}|\theta)] \\
&= \arg \max_{\mathbf{Q}} -\frac{1}{2} \mathbb{E} \left[ \mathbf{u}^T \tilde{\mathbf{Q}}^{-1} \mathbf{u} \right] - \frac{D-1}{2} \log |\mathbf{Q}| \\
&\quad - \frac{1}{2} \mathbb{E} \left[ \mathbf{u}^T \sum_{i=1}^P \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{C}}_i^T \mathbf{u} \right] + \mathbb{E} [\mathbf{u}^T \sum_{i=1}^P \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{h}}_i] \\
&\quad - \frac{1}{2} \sum_{i=1}^P \tilde{\mathbf{h}}_i^T \mathbf{A} \tilde{\mathbf{h}}_i + \sum_{i=1}^P \log f_i(\gamma_i) + C_4
\end{aligned} \tag{18}$$

$$\begin{aligned}
\mathbf{Q}^{t+1} &= \arg \max_{\mathbf{Q}} -\frac{1}{2} \text{Tr} \left( \tilde{\mathbf{Q}}^{-1} (\Phi + \phi \phi^T) \right) \\
&\quad - \frac{D-1}{2} \log |\mathbf{Q}| \\
&= \arg \max_{\mathbf{Q}} -\frac{1}{2} \sum_{d=1}^{D-1} \text{Tr} \left( \mathbf{Q}^{-1} (\Phi_d + \phi_d \phi_d^T) \right) \\
&\quad - \frac{D-1}{2} \log |\mathbf{Q}| \tag{19}
\end{aligned}$$

where  $C_4$  is a constant,  $\phi_d$  is the  $d$ th vector of size  $K$  in  $\phi$ , and  $\Phi_d$  is the  $d$ th matrix of size  $K \times K$  on the diagonal of  $\Phi$ . The last equality follows from the fact that  $\tilde{\mathbf{Q}}$  is block diagonal. We used (6) for  $\mathbb{E} [\mathbf{u} \tilde{\mathbf{Q}}^{-1} \mathbf{u}]$ . Similar to (8) and (12) we find

$$\mathbf{Q}^{t+1} = \frac{1}{D-1} \sum_{d=1}^{D-1} (\Phi_d^t + \phi_d^t (\phi_d^t)^T). \tag{20}$$

Since (15) holds with equality for  $\gamma_i = \eta_i$ , and the curvature does not depend on  $\eta_i$ , it can be shown that the optimal value for  $\gamma_i$  is  $\tilde{\mathbf{C}}_i^T \phi$  [5]. Note that using the factor scores  $\mathbf{u}_d \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$  the mean vector  $\beta$  in (3) is not needed.

### 3.7. Factor Loadings

Finally, combining (5), (10) and (18) we compute  $\mathbf{c}_i^{t+1}$  as follows

$$\begin{aligned}
\mathbf{c}_i^{t+1} &= \arg \max_{\mathbf{c}_i} \mathbb{E} [\{\mathbf{y}_n, \mathbf{x}_n, \mathbf{z}_m, \mathbf{w}_m, \mathbf{h}_i, \mathbf{u}\} | \theta] \\
&= \arg \max_{\mathbf{c}_i} \sum_{n=1}^N \left( -\mathbb{E} \left[ e^{\mathbf{c}_i^T \mathbf{x}_n} \right] + y_{in} \mathbf{c}_i^T \boldsymbol{\xi}_n \right) \\
&\quad - \frac{M}{2\sigma_i^2} \mathbf{c}_i^T \mathbf{B} \mathbf{c}_i - \frac{1}{2\sigma_i^2} \sum_{m=1}^M (\mathbf{c}_i^T \mathbf{a}_m - z_{im})^2 \\
&\quad - \frac{1}{2} \text{Tr} \left( \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{C}}_i^T (\Phi + \phi \phi^T) \right) + \phi^T \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{h}}_i. \tag{21}
\end{aligned}$$

Note that  $\mathbf{c}_i^T \mathbf{x}_n | \{\mathbf{y}_n\} \sim \mathcal{N}(\mathbf{c}_i^T \boldsymbol{\xi}_n, \mathbf{c}_i^T \Psi_n \mathbf{c}_i)$ . Completing the square in the Gaussian integral it is straightforward to show

that  $\mathbb{E} \left[ e^{\mathbf{c}_i^T \mathbf{x}_n} \right] = e^{\mathbf{c}_i^T \boldsymbol{\xi}_n + \frac{\mathbf{c}_i^T \Psi_n \mathbf{c}_i}{2}}$ . Moreover, we can write  $\text{Tr} \left( \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{C}}_i^T (\Phi + \phi \phi^T) \right) = \mathbf{c}_i^T \mathbf{U} \mathbf{c}_i$ , and  $\phi^T \tilde{\mathbf{C}}_i \mathbf{A} \tilde{\mathbf{h}}_i = \mathbf{c}_i^T \boldsymbol{\delta}_i$ , where  $\mathbf{U} = \frac{D-1}{2} \mathbf{Q}^{t+1} - \mathbf{V}$ ,  $\mathbf{V} = [\psi_1 \cdots \psi_K]$ ,  $\psi_k = \frac{1}{2D} \sum_{\ell=0}^{D-2} \sum_{d=1}^{D-1}$  and  $\mathbf{F} = [\phi_1 \cdots \phi_{D-1}]$ .

We can use Newton's method to find  $\mathbf{c}_i^{t+1}$  through the gradient and the Hessian, i.e.,

$$\begin{aligned}
\nabla_{\mathbf{c}_i} &= - \sum_{n=1}^N e^{\mathbf{c}_i^T \boldsymbol{\xi}_n + \frac{\mathbf{c}_i^T \Psi_n \mathbf{c}_i}{2}} (\boldsymbol{\xi}_n + \Psi_n \mathbf{c}_i) + \sum_{n=1}^N y_{in} \boldsymbol{\xi}_n \\
&\quad - \frac{M}{\sigma_i^2} \mathbf{B} \mathbf{c}_i - \frac{1}{\sigma_i^2} \sum_{m=1}^M (\mathbf{c}_i^T \mathbf{a}_m - z_{im}) \mathbf{a}_m \\
&\quad - \sum_{d=1}^{D-1} a_d (\Phi_d + \phi_d \phi_d^T) \mathbf{c}_i + \mathbf{F} \mathbf{A} \tilde{\mathbf{h}}_i \\
\nabla_{\mathbf{c}_i}^2 &= - \sum_{n=1}^N e^{\mathbf{c}_i^T \boldsymbol{\xi}_n + \frac{\mathbf{c}_i^T \Psi_n \mathbf{c}_i}{2}} [(\boldsymbol{\xi}_n + \Psi_n \mathbf{c}_i)(\boldsymbol{\xi}_n + \Psi_n \mathbf{c}_i)^T + \Psi_n] \\
&\quad - \frac{1}{\sigma_i^2} \sum_{m=1}^M \mathbf{a}_m \mathbf{a}_m^T - \sum_{d=1}^{D-1} a_d (\Phi_d + \phi_d \phi_d^T). \tag{22}
\end{aligned}$$

## 4. EXPERIMENTS

We test our algorithm on a Twitter dataset that is filtered from the 10% of the tweets in January 2013. In our dataset, we consider 2444 hashtags as objects (i.e.,  $P = 2444$ ); and analyze their number of occurrences in 743 hours (i.e.,  $N = 743$ ), available geographical coordinates, and word counts. We model the hashtag occurrences (i.e., number of tweets that mention a hashtag) using Poisson distribution. Word counts are modeled using multinomial distribution with a dictionary size of 2645 after eliminating the words that appear less than 100 times in the whole dataset.

The number of available coordinates ranges between 10 and 10456 with a mean of 134. Since geographical coordinates (latitude and longitude) constitute spherical data, they are better modeled using von Mises-Fisher (vMF) distribution than Gaussian, which is treated initially due to its popularity. Thus, we here present an extension of our algorithm for vMF distribution. The observation model for vMF distribution is given below

$$\begin{aligned}
\mathbf{z}_{im}^T | \mathbf{W}_m &\sim \text{vMF}(\mathbf{c}_i^T \mathbf{W}_m, \kappa_i), \\
\mathbf{w}_{mk}^T &\sim \text{vMF}(\boldsymbol{\alpha}_k, s_k),
\end{aligned}$$

for  $m = 1, \dots, M_i$ , where  $\mathbf{z}_{im}^{3 \times 1}$  is the spherical coordinate vector obtained from the original latitude and longitude information;

$\mathbf{W}_m^{K \times 3} = [\mathbf{w}_{m1} \cdots \mathbf{w}_{mK}]^T$  is the latent factor scores for all three dimensions. The scores  $\mathbf{w}_{mk}^{3 \times 1}$  for each factor  $k$  are also vMF distributed.

The vMF likelihood is given by

$$P(\mathbf{z}_{im} | \mathbf{W}_m, \theta) = C(\kappa_i) e^{\kappa_i \mathbf{W}_m^T \mathbf{c}_i \mathbf{z}_{im}},$$

where  $C(\kappa_i) = \frac{\kappa_i}{2\pi(e^{\kappa_i} - e^{-\kappa_i})}$ ,  $\kappa_i$  is the concentration parameter, and  $\mathbf{c}_i^T \mathbf{W}_m$  is the mean direction. Combining the likelihood with the prior it is straightforward to show that the posterior  $P(\mathbf{w}_{mk} | \mathbf{z}_{im}, \theta)$  is also vMF with mean and concentration given by

$$\mathbf{a}_{mk} = \frac{s_k \boldsymbol{\alpha}_k + \kappa_i \mathbf{c}_i \mathbf{z}_{im}}{b_{mk}}, \quad b_{mk} = \|s_k \boldsymbol{\alpha}_k + \kappa_i \mathbf{c}_i \mathbf{z}_{im}\|,$$

respectively. Maximizing the expected complete-data log-likelihood  $E[\log P(\{\mathbf{z}_{im}, \mathbf{w}_{mk}\} | \theta)]$  we get

$$\boldsymbol{\alpha}_k^{t+1} = \frac{\sum_{m=1}^M \mathbf{a}_{mk}^t}{\|\sum_{m=1}^M \mathbf{a}_{mk}^t\|}, \quad M = \max(M_i).$$

Using the well-known approximation we estimate the concentration parameters as

$$s_k^{t+1} = \frac{3\bar{r}_k - \bar{r}_k^3}{1 - \bar{r}_k^2}, \quad \bar{r}_k = \frac{\|\sum_{m=1}^M \mathbf{a}_{mk}^t\|}{M}$$

$$\kappa_i = \frac{3\bar{R}_i - \bar{R}_i^3}{1 - \bar{R}_i^2}, \quad \bar{R}_i = \frac{\|\sum_{m=1}^{M_i} \mathbf{z}_{im}\|}{M_i}.$$

We choose the number of factors through experiments based on the perplexity performance for the Poisson observations.

Table 1 lists the dominant hashtags in five selected factors. Our algorithm localizes factors in three aspects: popularity, geography, and topic. On the other hand, a unimodal factor analyzer handles only one of those. For example, the first factor in Table 1 is about European soccer, and the dominant hashtags in this factor are all popular. Average number of occurrences over 743 hours are given in parentheses. Similarly hashtags related to the second factor share the common topic of new year, and are on average tweeted from US. Third factor is about Middle East countries. The reason Bahrain was so popular in January 2013 seems to be the Gulf Cup of Nations, a soccer tournament organized in Bahrain between

**Table 1.** Factors localized in terms of popularity, geography, and topic. Numbers in parentheses denote the average popularity scores.

Europe Soccer	US New Year	Middle East Countries	Astrology	East Coast Sports
#FCF (83)	#HappyNewYear (49)	#bahrain (133)	#Aries (210)	#Patriots (29)
#FACup (40)	#NewYear (13)	#Pakistan (30)	#Pisces (231)	#Knicks (18)
#Arsenal (33)	#newyearseve (3)	#India (11)	#Capricorn (210)	#GoHawks (20)
#ballondor (17)	#Best2012Memories (4)	#kuwait (5)	#Virgo (192)	#Steelers (2)
#realmadrid (14)	#Feliz2013 (1)	#Iraq (2)	#Sagittarius (216)	#Yankees (3)

January 5th and 18th. Astrology hashtags have regularly high counts. The last factor focuses on the East Coast sports teams in football, basketball, and baseball.

## 5. CONCLUSION

A generative graphical model and an EM algorithm to analyze it have been proposed to be able summarize multimodal objects that consist of disparate observations. In a one-month Twitter dataset, the discovered factors, each of which acts as a summary multimodal object, have been shown to successfully localize hashtags in terms of popularity, geography, and topic.

## 6. REFERENCES

- [1] Mingyuan Zhou, Lauren A. Hannah, David B. Dunson, and Lawrence Carin, “Beta-negative binomial process and poisson factor analysis,” in *AISTATS*, 2012.
- [2] David M. Blei, Andrew Y. Ng, and Michael I. Jordan, “Latent dirichlet allocation,” *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.
- [3] J. Ngiam, A. Khosla, B. Kim, J. Nam, H. Lee, and A.Y. Ng, “Multimodal deep learning,” in *ICML*, 2011.
- [4] Y. Yılmaz, G.V. Moustakides, and X. Wang, “Sequential joint detection and estimation,” *Theory of Probability & Its Applications*, vol. 59, no. 3, 2014.
- [5] M.E. Khan, G. Bouchard, B.M. Marlin, and K.P. Murphy, “Variational bounds for mixed-data factor analysis,” in *Neural Information Processing Systems (NIPS) Conference*, 2010.
- [6] A.C. Smith and E.N. Brown, “Estimating a state-space model from point process observations,” *Neural Comput.*, vol. 15, no. 5, pp. 965–991, 2003.
- [7] D Böhning, “Multinomial logistic regression algorithm,” *Ann. Inst. Statist. Math.*, vol. 44, no. 1, pp. 197–200, 1992.