

Supplementary: Topic-Based Embeddings for Learning from Large Knowledge Graphs

1 Gibbs Sampling for BPBFM-2

Just as we did for BPBFM-1, we can express each latent count \mathcal{X}_{ij}^r in BPBFM-2 (which models each Λ^r as $\Lambda^r = \sum_{m=1}^M \eta_{mr} G^m$) as a sum of the following form: $\mathcal{X}_{ij}^r = \sum_{k_1=1}^K \sum_{k_2=1}^K \sum_{m=1}^M X_{ik_1 k_2 m j}^r$ where $\mathcal{X}_{ik_1 k_2 m j}^r \sim \text{Poisson}(u_{ik_1} \eta_{mr} G_{k_1 k_2}^m u_{jk_2})$. We further define $\mathcal{X}_{\cdot, k_1 k_2 m \cdot}^r = \sum_i \sum_j \sum_{r=1}^R \mathcal{X}_{ik_1 k_2 m j}^r$ and $\mathcal{X}_{\dots m}^r = \sum_i \sum_j \sum_{k_1=1}^K \sum_{k_2=1}^K \mathcal{X}_{ik_1 k_2 m j}^r$, using additive property of Poisson distribution

$$\mathcal{X}_{\cdot, k_1 k_2 m \cdot}^r \sim \text{Poisson}(\theta_{k_1 k_2} G_{k_1 k_2}^m \sum_{r=1}^R \eta_{mr}) \quad (1)$$

$$\mathcal{X}_{\dots m}^r \sim \text{Poisson}(\eta_{mr} \sum_{k_1=1}^K \sum_{k_2=1}^K \theta_{k_1 k_2} G_{k_1 k_2}^m) \quad (2)$$

With these defined, we proceed to give the update equations for the Gibbs sampler for BPBFM-2.

Sampling \mathcal{X}_{ij}^r : \mathcal{X}_{ij}^r is sampled just as in model-1.

Sampling $\mathcal{X}_{ik_1 k_2 m j}^r$: $\mathcal{X}_{ik_1 k_2 m j}^r$ can be sampled as

$$\mathcal{X}_{ik_1 k_2 j}^r \sim \text{Mult}(\mathcal{X}_{ij}^r; \frac{u_{ik_1} \eta_{mr} G_{k_1 k_2}^m u_{jk_2}}{\sum_{k_1=1}^K \sum_{k_2=1}^K u_{ik_1} \Lambda_{k_1 k_2}^r u_{jk_2}}) \quad (3)$$

Sampling $\mathbf{U}_{:,k}$: Using Dirichlet-multinomial conjugacy, each column of \mathbf{U} can be sampled as

$$\mathbf{U}_{:,k} \sim \text{Dir}(a + \mathcal{X}_{1k \dots}, a + \mathcal{X}_{2k \dots}, \dots, a + \mathcal{X}_{Nk \dots}) \quad (4)$$

where $\mathcal{X}_{ik \dots} = \sum_{k_2=1}^K \sum_{j=1}^N \sum_{m=1}^M \sum_{r=1}^R \mathcal{X}_{ik_2 m j}^r$.

Sampling $d_{k_1 k_2}^m$: Marginalizing out $G_{k_1 k_2}^m$ from Eq.(1), we have

$$\mathcal{X}_{\cdot, k_1 k_2 m \cdot}^r \sim \text{NegBin}((\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}}, p_{k_1 k_2}) \quad (5)$$

where $p_{k_1 k_2} = \frac{\theta_{k_1 k_2} \sum_{r=1}^R \eta_{mr}}{\theta_{k_1 k_2} \sum_{r=1}^R \eta_{mr} + \beta}$. Using the data augmentation scheme proposed we used for BPBFM-1, $d_{k_1 k_2}^m$ can be sampled by first sampling

$$\ell_{kk_2}^m \sim \sum_{t=1}^{X_{\cdot, k_1 k_2 m \cdot}^r} \text{Bern}\left(\frac{(\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}}}{(\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}} + t - 1}\right) \quad (6)$$

and then sampling

$$d_{k_1 k_2}^m \sim \text{Ga}\left(\frac{\gamma_0}{K} + \sum_{k_2} \ell_{kk_2}^m, \frac{1}{c_0 - \sum_{k_2} (\epsilon^m)^{\delta_{k_1 k_2}} (d_{k_2}^m)^{1-\delta_{k_1 k_2}} \ln(1 - p_{k_1 k_2})}\right)$$

Sampling ϵ^m : ϵ^m can be sampled as

$$\epsilon^r \sim \text{Ga}\left(e_0 + \sum_k \ell_{kk}^r, \frac{1}{f_0 - \sum_k d_k^m \ln(1 - p_{kk})}\right) \quad (8)$$

Sampling $G_{k_1 k_2}^m$: Using Gamma-Poisson conjugacy, $G_{k_1 k_2}^m$ can be sampled by

$$G_{k_1 k_2}^m \sim \text{Ga}((\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^r (d_{k_2}^r)^{1-\delta_{k_1 k_2}} + X_{\cdot, k_1 k_2 m \cdot}^r, \frac{1}{\beta + \theta_{k_1 k_2} \sum_{r=1}^R \eta_{mr}}) \quad (9)$$

Sampling η_{mr} : Using equation (2) and Gamma-Poisson conjugacy, η_{mr} can be sampled by

$$\eta_{mr} \sim \text{Ga}(h_0 + \mathcal{X}_{\dots m}^r, \frac{1}{q_0 + \sum_{k_1=1}^K \sum_{k_2=1}^K \theta_{k_1 k_2} G_{k_1 k_2}^m}) \quad (10)$$

2 Online Gibbs Sampling

In this section, we provide the details of the online Gibbs sampling algorithms for both of our models. Our online Gibbs sampling algorithms are based on the idea of the recently developed Bayesian Conditional Density Filtering (BCDF) framework (Guhaniyogi et al., 2014). Their key idea in BCDF is to process data in small minibatches, and maintain and update sufficient statistics of the model parameters with each new minibatch of the data. In our models, these sufficient statistics are the latent counts.

2.1 Online Gibbs Sampling for BPBFM-1

Denoting I_t as indices of valid triplets in minibatch selected at iteration t , and I as the indices of all the valid triplets in training data. Define $\mathcal{X}_{ik \dots}^{r,t} = \frac{|I|}{|I_t|} \sum_{k_2=1}^K \sum_{j=1, ij \in I_t}^N \mathcal{X}_{ik_2 j}^r$, $\mathcal{X}_{ik \dots}^t = \frac{|I|}{|I_t|} \sum_{r=1}^R \mathcal{X}_{ik \dots}^{r,t}$, and $\mathcal{X}_{\cdot, k_1 k_2 \cdot}^{r,t} = \frac{|I|}{|I_t|} \sum_{i,j \in I_t} \mathcal{X}_{ik_1 k_2 j}^r$, where $|I|$ and $|I_t|$ are cardinalities of the two sets. Then similar to batch Gibbs Sampling, define following quantities for $t \leq 2$:

$$\begin{aligned} \mathcal{X}_{ik \dots}^{r,t} &= (1 - \rho) \mathcal{X}_{ik \dots}^{r,t-1} + \rho \frac{|I|}{|I_t|} \sum_{k_2=1}^K \sum_{j=1, ij \in I_t}^N \mathcal{X}_{ik_2 j}^r, \\ \mathcal{X}_{ik \dots}^t &= (1 - \rho) \mathcal{X}_{ik \dots}^{t-1} + \rho \frac{|I|}{|I_t|} \sum_{r=1}^R \mathcal{X}_{ik \dots}^{r,t}, \text{ and} \\ \mathcal{X}_{\cdot, k_1 k_2 \cdot}^{r,t} &= (1 - \rho) \mathcal{X}_{\cdot, k_1 k_2 \cdot}^{r,t-1} + \rho \frac{|I|}{|I_t|} \sum_{i,j \in I_t} \mathcal{X}_{ik_1 k_2 j}^r. \end{aligned} \quad (7)$$

Here $\rho = (t + t_0)^{-w}$ is a decaying learning rate,

as used in other online inference algorithms, such as stochastic variational inference (Hoffman et al., 2013). Here, $t_0 > 0$ and $w \in (0.5, 1]$ are required to guarantee convergence. With these defined, online Gibbs sampling at iteration t proceeds as:

Sampling $\mathbf{U}_{:,k}$: Each column of \mathbf{U} can be sampled as

$$\mathbf{U}_{:,k} \sim \text{Dir}(a + \mathcal{X}_{1k\dots}^t, a + \mathcal{X}_{2k\dots}^t, \dots, a + \mathcal{X}_{Nk\dots}^t) \quad (11)$$

Sampling d_k^r : d_k^r can be sampled by first sampling

$$\ell_{kk_2}^r \sim \sum_{t=1}^{\mathcal{X}_{\cdot k_1 k_2}^{r,t}} \text{Bern}\left(\frac{(\epsilon^r)^{\delta_{k_1 k_2}} d_{k_1}^r (d_{k_2}^r)^{1-\delta_{k_1 k_2}}}{(\epsilon^r)^{\delta_{k_1 k_2}} d_{k_1}^r (d_{k_2}^r)^{1-\delta_{k_1 k_2}} + t - 1}\right) \quad (12)$$

and then sampling

$$d_k^r \sim \text{Ga}\left(\frac{\gamma_0}{K} + \sum_{k_2} \ell_{kk_2}^r, \frac{1}{c_0 - \sum_{k_2}^K (\epsilon^r)^{\delta_{kk_2}} (d_{k_2}^r)^{1-\delta_{kk_2}} \ln(1 - p_{kk_2})}\right) \quad (13)$$

Sampling ϵ^r : ϵ^r can be sampled as

$$\epsilon^r \sim \text{Ga}\left(e_0 + \sum_k \ell_{kk}^r, \frac{1}{f_0 - \sum_k d_k^r \ln(1 - p_{kk})}\right) \quad (14)$$

Sampling $\Lambda_{k_1 k_2}^r$: $\Lambda_{k_1 k_2}^r$ can be sampled by

$$\Lambda_{k_1 k_2}^r \sim \text{Ga}\left((\epsilon^r)^{\delta_{k_1 k_2}} d_{k_1}^r (d_{k_2}^r)^{1-\delta_{k_1 k_2}} + \mathcal{X}_{\cdot k_1 k_2}^{r,t}, \frac{1}{\beta + \theta_{k_1 k_2}}\right) \quad (15)$$

\mathcal{X}_{ij}^r , ϵ^r and $\mathcal{X}_{ik_1 k_2 j}^r$ are sampled the same way as the batch Gibbs sampling.

2.2 Online Gibbs Sampling for BPBFM-2

Similar to online BPBFM-1, we define $\mathcal{X}_{\cdot k_1 k_2 m}^{r,t} = (1-\rho)\mathcal{X}_{\cdot k_1 k_2 m}^{r,t-1} + \rho \frac{|I|}{|I_t|} \sum_{ij, ij \in I_t} \sum_{r=1}^R \mathcal{X}_{ik_1 k_2 m j}^r$, $\mathcal{X}_{\dots m}^{r,t} = (1-\rho)\mathcal{X}_{\dots m}^{r,t-1} + \rho \frac{|I|}{|I_t|} \sum_{ij, ij \in I_t} \sum_{k_1=1}^K \sum_{k_2=1}^K \mathcal{X}_{ik_1 k_2 m j}^r$, and $\mathcal{X}_{ik\dots}^{r,t} = (1-\rho)\mathcal{X}_{ik\dots}^{r,t-1} + \rho \frac{|I|}{|I_t|} \sum_{k_2=1}^K \sum_{j=1, ij \in I_t}^N \sum_{m=1}^M \sum_{r=1}^R \mathcal{X}_{ik k_2 m j}^r$. With these defined, we proceed to give the update equations for the online Gibbs sampler for model-2:

Sampling $\mathbf{U}_{:,k}$: Each column of \mathbf{U} can be sampled as

$$\mathbf{U}_{:,k} \sim \text{Dir}(a + \mathcal{X}_{1k\dots}^t, a + \mathcal{X}_{2k\dots}^t, \dots, a + \mathcal{X}_{Nk\dots}^t) \quad (16)$$

Sampling d_m^r : d_m^r can be sampled by first sampling

$$\ell_{kk_2}^m \sim \sum_{t=1}^{\mathcal{X}_{\cdot k_1 k_2 m}^{r,t}} \text{Bern}\left(\frac{(\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}}}{(\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}} + t - 1}\right) \quad (17)$$

and then sampling

$$d_k^m \sim \text{Ga}\left(\frac{\gamma_0}{K} + \sum_{k_2} \ell_{kk_2}^m, \frac{1}{c_0 - \sum_{k_2}^K (\epsilon^m)^{\delta_{kk_2}} (d_{k_2}^m)^{1-\delta_{kk_2}} \ln(1 - p_{kk_2})}\right) \quad (18)$$

Sampling $G_{k_1 k_2}^m$: Using Gamma-Poisson conjugacy, $G_{k_1 k_2}^m$ can be sampled by

$$G_{k_1 k_2}^m \sim \text{Ga}\left((\epsilon^m)^{\delta_{k_1 k_2}} d_{k_1}^m (d_{k_2}^m)^{1-\delta_{k_1 k_2}} + \mathcal{X}_{\cdot k_1 k_2 m}^{r,t}, \frac{1}{\beta + \theta_{k_1 k_2} \sum_{r=1}^R \eta_{mr}}\right) \quad (19)$$

Sampling η_{mr} : η_{mr} can be sampled by

$$\eta_{mr} \sim \text{Ga}\left(h_0 + \mathcal{X}_{\dots m}^{r,t}, \frac{1}{q_0 + \sum_{k_1=1}^K \sum_{k_2=1}^K \theta_{k_1 k_2} G_{k_1 k_2}^m}\right) \quad (20)$$

\mathcal{X}_{ij}^r , $\mathcal{X}_{ik_1 k_2 m j}^r$, ϵ^m ϵ^m can be sampled the same way as the batch Gibbs sampling.

References

- R. Guhaniyogi, S. Qamar, and D. B. Dunson. Bayesian conditional density filtering. *arXiv preprint arXiv:1401.3632*, 2014.
- M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *JMLR*, 2013.