Probabilistic Collaborative Representation Learning for Personalized Item Recommendation
Supplementary Material

Aghiles Salah and **Hady W. Lauw**
School of Information Systems
Singapore Management University, Singapore
{asalah, hadywlauw}@smu.edu.sg

A VARIATIONAL PCRL

In this supplement we provide the details of the derivation of the variational PCRL algorithm.

A.1 UPDATES OF $\tilde{\boldsymbol{\lambda}}^{\theta}$ AND $\tilde{\boldsymbol{\phi}}$

Thanks to the auxiliary variables $\boldsymbol{s}$, the collaborative-specific part of PCRL is *conditionally conjugate* (Ghahramani and Beal, 2001): the complete conditional[1] of each θ_{uk} ($\boldsymbol{s}_{ui}$) and its corresponding variational distribution are in the same exponential family. Thereby, we can perform the updates of $\tilde{\boldsymbol{\lambda}}^{\theta}$ and $\tilde{\boldsymbol{\phi}}$ in closed-form.

The complete conditionals over the user preferences θ_{uk} and auxiliary variables $\boldsymbol{s}_{ui}$ are respectively:

$$\theta_{uk}|\mathbf{X}, \boldsymbol{\beta}, \boldsymbol{s} \sim \text{Gamma}(\lambda^s_\theta + \textstyle\sum_i s_{uik}, \lambda^r_\theta + \sum_i \beta_{ik})$$

$$\boldsymbol{s}_{ui}|\mathbf{X}, \boldsymbol{\theta}, \boldsymbol{\beta} \sim \text{Multinomial}(x_{ui}, \log \boldsymbol{\phi}_{ui}).$$

where $\boldsymbol{\phi}_{ui}$ is a point on the K-simplex, and for all k: $\phi_{uik} \propto \theta_{uk}\beta_{ik}$.

The optimal coordinate updates for the parameters $\tilde{\boldsymbol{\lambda}}^{\theta}$ and $\tilde{\boldsymbol{\phi}}$ are given by the expected natural parameters (w.r.t. q) of the above conditionals, which yields:

$$\tilde{\lambda}^{\theta}_{uk} = \left(\lambda^s_\theta + \textstyle\sum_i x_{ui}\tilde{\phi}_{uik}, \lambda^r_\theta + \sum_i \frac{\tilde{\lambda}^{\beta,s}_{ik}}{\tilde{\lambda}^{\beta,r}_{ik}}\right),$$

$$\tilde{\phi}_{uik} \propto \exp\left(\psi(\tilde{\lambda}^{\theta,s}_{uk}) - \log\tilde{\lambda}^{\theta,r}_{uk} + \psi(\tilde{\lambda}^{\beta,s}_{ik}) - \log\tilde{\lambda}^{\beta,r}_{ik}\right),$$

where we have used the standard results about the expectation of Gamma and Multinomial random variables. That is, if $\theta \sim \text{Gamma}(\lambda^s, \lambda^r)$, then $\mathbb{E}(\theta) = \frac{\lambda^s}{\lambda^r}$ and $\mathbb{E}(\log\theta) = \psi(\lambda^s) - \log\lambda^r$, with $\psi(\cdot)$ denoting the digamma function. If $\boldsymbol{s}_{ui} \sim \text{Multinomial}(x_{ui}, \boldsymbol{\phi}_{ui})$, then the expectation of the k^{th} component of $\boldsymbol{s}_{ui}$ is $\mathbb{E}(z_{uik}) = x_{ui}\phi_{uik}$.

[1]The conditional distribution of a variable given the other variables and observations

A.2 MONTE CARLO ESTIMATOR OF THE GRADIENT OF THE ELBO.

In this section, we first review the rejection sampling procedure and its reparameterized variant. We then, give the derivation details of equations (10) and (11).

A.2.1 Reparameterized Acceptance-Rejection Sampling

Rejection sampling (Robert and Casella, 2005) is a widely used technique to simulate random variables from complex distributions. To sample from $q(\beta;\omega)$ using this method, we first introduce a *proposal distribution* $r(\beta;\omega)$ that is easy to sample from, and which satisfies $q(\beta;\omega) \leq A_\omega r(\beta;\omega)$ for some constant $A_\omega < \infty$. Next, (i) we generate $\beta \sim r(\beta;\omega)$ and $u \sim \mathcal{U}[0,1]$, (ii) if $u > \frac{q(\beta;\omega)}{A_\lambda r(\beta;\omega)}$ then we reject β and return to i, (iii) we accept β, otherwise.

In the reparameterized variant of this procedure (Naesseth et al., 2017), we further require that sampling from the proposal distribution can be carried out using the reparameterization trick, i.e., $\beta \sim r(\beta;\omega)$ is equivalent to $\beta = \mathcal{G}(\epsilon;\omega)$ with $\epsilon \sim t(\epsilon)$. The accepted samples are then ϵ instead of β.

The key to use the reparameterized rejection sampling to form a Monte Carlo estimator of the gradient of the ELBO, is the marginal distribution of the accepted sample ϵ noted $\pi(\epsilon;\omega)$. Fortunately, this distribution is available and can be obtained by marginalizing out the uniform random variable u as follows

$$\begin{aligned}\pi(\epsilon;\omega) &= \int \pi(\epsilon, u; \lambda) du \\ &= \int A_\omega t(\epsilon) \mathbb{1}\left[0 < u < \frac{q(\mathcal{G}(\epsilon,\omega);\omega)}{A_\omega r(\mathcal{G}(\epsilon,\omega);\omega)}\right] du \\ &= t(\epsilon)\frac{q(\mathcal{G}(\epsilon,\omega);\omega)}{r(\mathcal{G}(\epsilon,\omega);\omega)},\end{aligned}$$

where $\mathbb{1}$ is the indicator function. The above marginal follows from the definition of the rejection sampling procedure, for more details please refer to (Naesseth et al., 2017).

A.2.2 EXPECTATION WITH RESPECT TO π AND DERIVATION OF THE GRADIENT

The first step in building the Monte Carlo Estimator is to rewrite $\mathbb{E}_{q(\boldsymbol{\beta}_i;\omega)}[\log p(\mathbf{c}_i|\mathcal{W},\boldsymbol{\beta}_i)]$ as an expectation with respect to $\pi(\boldsymbol{\epsilon}_i;\omega)$. The details of this step given below.

$$
\begin{aligned}
&\mathbb{E}_{q(\boldsymbol{\beta}_i;\omega)}[\log p(\mathbf{c}_i|\mathcal{W},\boldsymbol{\beta}_i)] \\
&= \int \frac{q(\boldsymbol{\beta}_i;\omega)}{r(\boldsymbol{\beta}_i;\omega)} r(\boldsymbol{\beta}_i;\omega)\log p(\mathbf{c}_i|\mathcal{W},\boldsymbol{\beta}_i)d\boldsymbol{\beta}_i \\
&= \int \frac{q(\mathcal{G}(\boldsymbol{\epsilon}_i,\omega);\omega)}{r(\mathcal{G}(\boldsymbol{\epsilon}_i,\omega);\omega)} t(\boldsymbol{\epsilon}_i)\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))d\boldsymbol{\epsilon}_i \\
&= \int \pi(\boldsymbol{\epsilon}_i;\omega)\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))d\boldsymbol{\epsilon}_i \\
&= \mathbb{E}_{\pi(\boldsymbol{\epsilon}_i;\omega)}[\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))]
\end{aligned}
$$

The second equality in the above equation holds since: $\beta = \mathcal{G}(\epsilon,\omega)$ with $\epsilon \sim t(\epsilon)$, is a reparameterization of $\beta \sim r(\beta;\omega)$.

The details of the computation of gradient (11) are as follows

$$
\begin{aligned}
&= \nabla_\omega \mathbb{E}_{\pi(\boldsymbol{\epsilon}_i;\omega)}[\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))] \\
&= \nabla_\omega \int \pi(\boldsymbol{\epsilon}_i;\omega)\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))d\boldsymbol{\epsilon}_i \\
&\overset{a}{=} \int \pi(\boldsymbol{\epsilon}_i;\omega)\nabla_\omega \log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))d\boldsymbol{\epsilon}_i \\
&+ \int \log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))\nabla_\omega \pi(\boldsymbol{\epsilon}_i;\omega)d\boldsymbol{\epsilon}_i \\
&\overset{b}{=} \int \pi(\boldsymbol{\epsilon}_i;\omega)\nabla_\omega \log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))d\boldsymbol{\epsilon}_i \\
&+ \int \pi(\boldsymbol{\epsilon}_i;\omega)\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))\nabla_\omega \log \pi(\boldsymbol{\epsilon}_i;\omega)d\boldsymbol{\epsilon}_i \\
&\overset{c}{=} \mathbb{E}_{\pi(\boldsymbol{\epsilon}_i;\omega)}[\nabla_\omega \log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))] \\
&+ \mathbb{E}_{\pi(\boldsymbol{\epsilon}_i;\omega)}[\log p(\mathbf{c}_i|\mathcal{W},\mathcal{G}(\boldsymbol{\epsilon}_i,\omega))\nabla_\omega \log \pi(\boldsymbol{\epsilon}_i;\omega)]
\end{aligned}
$$

where we pushed the gradient into the integral, assuming the necessary regularity conditions, and used the derivative rule of a product in a. In b we made use the log derivative-trick (also known as REINFORCE) (Glynn, 1990; Williams, 1992), i.e., $\nabla q = q\nabla \log q$. Finally, we rewrote integrals as expectations in c.

B Category Labels

The ten top categories retained in each dataset to evaluate the quality of item representations are listed below.

Office: “Writing & Correction Supplies”, “Paper”, “Telephones & Accessories”, “Printer Ink & Toner”, “Desk Accessories & Workspace Organizers”, “Printers & Accessories”, “Labels, Indexes & Stamps”, “Binders & Binding Systems”, “Tape, Adhesives & Fasteners”, “Envelopes, Mailers & Shipping Supplies”.

Automotive: “Bulbs”, “Protective Gear”, “Filters”, “Towing Products & Winches”, “Exterior Care”, “Lighting & Electrical”, “Accessories”, “Parts”, “Shocks Struts & Suspension”, “Brake System”.

Sports: “Fishing”, “Hunting”, “Accessories”, “Camping & Hiking”, “Archery”, “Tactical & Duty”, “Airsoft”, “Men”, “Strength Training Equipment”, “Parts & Components”.

Pet Supplies: “Grooming”, “Toys”, “Health Supplies”, “Food”, “Feeding & Watering Supplies”, “Litter & Housebreaking”, “Treats”, “Beds & Furniture”, “Collars, Harnesses & Leashes”, “Apparel & Accessories”.

References

Ghahramani, Z. and Beal, M. J. (2001). Propagation algorithms for variational bayesian learning. In *NIPS*, pages 507–513.

Glynn, P. W. (1990). Likelihood ratio gradient estimation for stochastic systems. *CACM*, 33(10):75–84.

Naesseth, C., Ruiz, F., Linderman, S., and Blei, D. (2017). Reparameterization gradients through acceptance-rejection sampling algorithms. In *AISTATS*, pages 489–498.

Robert, C. and Casella, G. (2005). *Monte Carlo Statistical Methods*. Springer New York.

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.