$$\mathcal{L}_{data} \triangleq \mathbb{E}_q \left[\log p(x \mid \mathcal{Z}, \phi) + \log \frac{p(\phi \mid \tau_1, \tau_0)}{q(\phi)} \right] \tag{16}$$

$$\mathcal{L}_{HDP} \triangleq \mathbb{E}_q \left[\log \frac{p(u \mid \gamma)}{q(u)} + \log \frac{p(\pi \mid u, \alpha)}{q(\pi)} \right] \tag{17}$$

$$\mathcal{L}_{alloc} \triangleq \mathbb{E}_q[\log p(\mathcal{Z} \mid \pi)] \tag{18}$$

and $\mathbb{H}[q(\mathcal{Z})]$ as the entropy of $q(\mathcal{Z})$. Here, we discuss the non-model-specific term $\mathcal{L}_{HDP}$. While most of its evaluation requires standard expectations of beta and Dirichlet distributions, difficulty arises in the term $\mathbb{E}_q[\log p(\pi)]$. The issue comes from the log-normalization constant, $c_D(\alpha\beta)$, of the log-Dirichlet distribution $\log p(\pi)$, where:

$$c_D(\alpha\beta) \triangleq \log \Gamma \left(\sum_{m=1}^{K+1} \alpha\beta_m \right) - \sum_{m=1}^{K+1} \log \Gamma(\alpha\beta_m) \tag{19}$$

In particular, $\mathbb{E}_q[c_D(\alpha\beta)]$, has no closed form due to our choice of $q(u)$. We place a lower bound on this term:

$$c_D(\alpha\beta) \geq K \log \alpha + \sum_{m=1}^{K+1} \log \beta_m, \tag{20}$$

a proof of which is given in Appendix A. The expectation of Eq. 20 has a closed form, giving us a tractable lower bound on $\mathcal{L}$.

5.2 Variational Inference for the HDP-HMM

Factor $q(z)$. We assume a form of $q(z)$ that retains the Markov dependencies from $p(z \mid \pi)$:

$$q(z) = q(z_1) \prod_{t=2}^{T} q(z_t \mid z_{t-1}) = \left[\prod_{\ell=1}^{K} \hat{r}_{1\ell}^{z_{1\ell}} \right] \left[\prod_{t=2}^{T} \prod_{\ell=1}^{K} \prod_{m=1}^{K} \left(\frac{\hat{s}_{t\ell m}}{\hat{r}_{t-1,\ell}} \right)^{z_{t-1,\ell} z_{tm}} \right], \tag{21}$$

where variational parameters $\hat{s}, \hat{r}$ parameterize the discrete distribution $q(z)$ by $\hat{s}_{t\ell m} \triangleq q(z_{t-1,\ell}, z_{tm})$ and $\hat{r}_{t\ell} \triangleq q(z_{t\ell})$. Note that, due to our truncation assumption, products run up to K rather than ∞. Following previous work on variational methods for HMMs [3], our update to $q(z)$ finds the joint configuration of all parameters $\hat{s}, \hat{r}$ that maximize our objective $\mathcal{L}$. To do so, we apply Eq. 11 to find that the optimal $q(z)$ satisfies:

$$q(z) \propto \left[\prod_{\ell=1}^{K} \tilde{\pi}_{0\ell}^{z_{1\ell}} \right] \left[\prod_{t=2}^{T} \prod_{\ell=1}^{K} \prod_{m=1}^{K} \tilde{\pi}_{\ell m}^{z_{t-1,\ell} z_{tm}} \right] \left[\prod_{t=1}^{T} \prod_{\ell=1}^{K} \exp\left(\mathbb{E}_q[\log p(x_t \mid \phi_\ell)] \right)^{z_{t\ell}} \right] \tag{22}$$

where $\tilde{\pi}_{\ell m}$ is as defined in Eq. 13. We can now use belief propagation – or more specifically, the forward-backward algorithm – to find the marginals $\hat{s}_{t\ell m} = q(z_{t-t,\ell}, z_{tm})$ and $\hat{r}_{t\ell} = q(z_{t\ell})$ of the distribution $q(z)$ satisfying this proportionality.

Factor $q(\pi)$. In the case of the HDP-HMM, $q(\pi)$ is a factor over an infinite number of distributions: $q(\pi) = \prod_{\ell=0}^{\infty} \text{Dirichlet}(\hat{\theta}_\ell)$; however, as noted in Section 5.1, our chosen truncation of $q(z)$ implies we need only update up to $\ell = K$, as terms with index $\ell > K$ will be set to their prior.

Following the generic update to $\hat{\theta}$ given in Eq. 14, we find that the usage sufficient statistics are given by:

$$U_{\ell m} = \begin{cases} \hat{r}_{1m} & \ell = 0 \\ \sum_{t=2}^{T} \hat{s}_{t\ell m} & \ell \neq 0 \end{cases} \tag{23}$$