
Appendix for “VAE Learning via Stein Variational Gradient Descent”

Yuchen Pu, Zhe Gan, Ricardo Henao, Chunyuan Li, Shaobo Han, Lawrence Carin

Department of Electrical and Computer Engineering, Duke University

{yp42, zg27, r.henao, cl319, shaobo.han, lcarin}@duke.edu

A Proof

Proof of Theorem 1 Recall the definition of KL divergence:

$$\text{KL}(q_T \| p) = \text{KL}(q_T(\boldsymbol{\theta})q(\mathcal{Z}) || p(\boldsymbol{\theta}, \mathcal{Z} | \mathcal{D})) = \int \int q_T(\boldsymbol{\theta})q(\mathcal{Z}) \log \frac{p(\boldsymbol{\theta}, \mathcal{Z} | \mathcal{D})}{q_T(\boldsymbol{\theta})q(\mathcal{Z})} d\boldsymbol{\theta} d\mathcal{Z} \quad (1)$$

$$= \int q_T(\boldsymbol{\theta}) \left\{ \int q(\mathcal{Z}) \log p(\boldsymbol{\theta}, \mathcal{Z}, \mathcal{D}) d\mathcal{Z} \right\} d\boldsymbol{\theta} \\ - \int q_T(\boldsymbol{\theta}) \log q_T(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int q(\mathcal{Z}) \log q(\mathcal{Z}) d\mathcal{Z} - \log p(\mathcal{D}) \quad (2)$$

$$= \int q_T(\boldsymbol{\theta}) \log \tilde{p}(\boldsymbol{\theta}; \mathcal{D}) d\boldsymbol{\theta} - \int q_T(\boldsymbol{\theta}) \log q_T(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int q(\mathcal{Z}) \log q(\mathcal{Z}) d\mathcal{Z} - \log p(\mathcal{D}) \quad (3)$$

$$= \text{KL}\left(q_T(\boldsymbol{\theta}) || \tilde{p}(\boldsymbol{\theta}; \mathcal{D})\right) - \int q(\mathcal{Z}) \log q(\mathcal{Z}) d\mathcal{Z} - \log p(\mathcal{D}), \quad (4)$$

where $\log \tilde{p}(\boldsymbol{\theta}; \mathcal{D}) = \int q(\mathcal{Z}) \log p(\boldsymbol{\theta}, \mathcal{Z}, \mathcal{D}) d\mathcal{Z}$. Since $\nabla_{\epsilon} \int q(\mathcal{Z}) \log q(\mathcal{Z}) d\mathcal{Z} = \nabla_{\epsilon_1} \log p(\mathcal{D}) = 0$, we have

$$\nabla_{\epsilon} \text{KL}(q_T(\boldsymbol{\theta})q(\mathcal{Z}) || p(\boldsymbol{\theta}, \mathcal{Z} | \mathcal{D})) = \nabla_{\epsilon} \text{KL}(q_T(\boldsymbol{\theta}) || \tilde{p}(\boldsymbol{\theta}; \mathcal{D})). \quad (5)$$

Following [2], we have

$$\nabla_{\epsilon} (\text{KL}(q_T(\boldsymbol{\theta})q(\mathcal{Z}) || p(\boldsymbol{\theta}, \mathcal{Z} | \mathcal{D}))|_{\epsilon_1=0}) \\ = -\mathbb{E}_{\boldsymbol{\theta} \sim q(\boldsymbol{\theta})} [\nabla_{\boldsymbol{\theta}} \log \tilde{p}(\boldsymbol{\theta}; \mathcal{D})^T \psi(\boldsymbol{\theta}; \mathcal{D}) + \text{trace}(\nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}; \mathcal{D}))] \quad (6)$$

$$= -\mathbb{E}_{\boldsymbol{\theta} \sim q(\boldsymbol{\theta})} [\text{trace}(\nabla_{\boldsymbol{\theta}} \log \tilde{p}(\boldsymbol{\theta}; \mathcal{D}) \psi(\boldsymbol{\theta}; \mathcal{D})^T + \psi(\boldsymbol{\theta}; \mathcal{D}))]. \quad (7)$$

Proof of Theorem 2 Following [1], we have $\mathbb{E}_{I=\{i_1, \dots, i_m\}} \left[\frac{1}{m} \sum_{i=j}^m a_{i_j} \right] = \frac{a_1 + \dots + a_k}{k}$, where $I \subset \{1, \dots, k\}$ with $|I| = m < k$, is a uniformly distributed subset of $\{1, \dots, k\}$. Using Jensen’s

inequality, we have

$$\text{KL}_{q,p}^k(\Theta; \mathcal{D}) = -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left[\log \frac{1}{k} \sum_{i=1}^k \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \quad (8)$$

$$= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left[\log \mathbb{E}_{I=\{i_1, \dots, i_m\}} \left[\frac{1}{m} \sum_{i=1}^m \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right] \quad (9)$$

$$\leq -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left[\mathbb{E}_{I=\{i_1, \dots, i_m\}} \left[\log \frac{1}{m} \sum_{i=1}^m \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right] \quad (10)$$

$$= -\mathbb{E}_{\Theta^{1:m} \sim q(\Theta)} \left[\log \frac{1}{m} \sum_{i=1}^m \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \quad (11)$$

$$= \text{KL}_{q,p}^m(\Theta; \mathcal{D}), \quad (12)$$

if $q(\Theta)/p(\Theta|\mathcal{D})$ is bounded, we have

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} = \mathbb{E}_{q(\Theta)} \left[\frac{p(\Theta | \mathcal{D})}{q(\Theta)} \right] = \int p(\Theta | \mathcal{D}) d\Theta = 1. \quad (13)$$

Therefore

$$\text{KL}_{q,p}^k(\Theta; \mathcal{D}) = -\lim_{k \rightarrow \infty} \mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left[\log 1 \right] = 0.$$

Proof of Theorem 3 $\mathcal{A}_p^k(\Theta^{1:k}; \mathcal{D})$ is defined as following:

$$\mathcal{A}_p^k(\Theta^{1:k}; \mathcal{D}) = \frac{1}{\tilde{\omega}} \sum_{i=1}^k \omega_i \left(\text{trace}(\mathcal{A}_p(\Theta^i; \mathcal{D})) \right) \quad (14)$$

$$\omega_i = p(\Theta^i; \mathcal{D})/q(\Theta^i), \quad \tilde{\omega} = \sum_{i=1}^k \omega_i \quad (15)$$

$$\mathcal{A}_p(\Theta; \mathcal{D}) = \nabla_{\Theta} \log \tilde{p}(\Theta; \mathcal{D}) \psi(\Theta; \mathcal{D})^T + \nabla_{\Theta} \psi(\Theta; \mathcal{D}). \quad (16)$$

Assume $p_{[T^{-1}]}(\Theta)$ denote the density of $\hat{\Theta} = T^{-1}(\Theta)$. We have

$$\nabla_{\epsilon} \left(\text{KL}_{q,p}^k(\Theta'; \mathcal{D}) \right) = -\nabla_{\epsilon} \left\{ \mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left[\log \frac{1}{k} \sum_{i=1}^k \frac{p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right\} \quad (17)$$

$$= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left\{ \nabla_{\epsilon} \left[\log \frac{1}{k} \sum_{i=1}^k \frac{p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right\} \quad (18)$$

$$= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left\{ \left[\frac{1}{k} \sum_{i=1}^k \frac{p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right]^{-1} \left[\frac{1}{k} \sum_{i=1}^k \frac{\nabla_{\epsilon} p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right\}. \quad (19)$$

Note that

$$\nabla_{\epsilon} p_{[T^{-1}]}(\Theta^i | \mathcal{D}) = p_{[T^{-1}]}(\Theta^i | \mathcal{D}) \nabla_{\epsilon} \log p_{[T^{-1}]}(\Theta^i | \mathcal{D}), \quad (20)$$

and when $\epsilon = 0$, we have

$$p_{[T^{-1}]}(\Theta^i | \mathcal{D}) = p(\Theta^i | \mathcal{D}), \quad \nabla_{\epsilon} T(\Theta) = \psi(\Theta; \mathcal{D}) \quad (21)$$

$$\nabla_{\epsilon} \nabla_{\Theta} T(\Theta) = \nabla_{\epsilon} \psi(\Theta; \mathcal{D}), \quad \nabla_{\Theta} T(\Theta) = \mathbf{I} \quad (22)$$

Therefore

$$\begin{aligned} \nabla_{\epsilon} \log p_{[T^{-1}]}(\Theta^i | \mathcal{D}) &= \nabla_{\epsilon} \log p(\Theta^i | \mathcal{D})^T \nabla_{\epsilon} T(\Theta^i) + \text{trace} \left((\nabla_{\Theta^i} T(\Theta^i))^{-1} \cdot \nabla_{\epsilon} \nabla_{\Theta^i} T(\Theta^i) \right) \\ &= \nabla_{\epsilon} \log p(\Theta^i | \mathcal{D})^T \psi(\Theta^i; \mathcal{D}) + \text{trace}(\nabla_{\epsilon} \psi(\Theta^i; \mathcal{D})) \end{aligned} \quad (23)$$

$$= \text{trace}(\nabla_{\epsilon} \log p(\Theta^i | \mathcal{D}) \psi(\Theta^i; \mathcal{D})^T + \nabla_{\epsilon} \psi(\Theta^i; \mathcal{D})) \quad (24)$$

$$= \text{trace}(\mathcal{A}_p(\Theta^i; \mathcal{D})). \quad (25)$$

Therefore, (19) can be rewritten as

$$\begin{aligned}
\nabla_{\epsilon} \left(\text{KL}_{q,p}^k(\Theta'; \mathcal{D}) \right) &= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left\{ \left[\frac{1}{k} \sum_{i=1}^k \frac{p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right]^{-1} \left[\frac{1}{k} \sum_{i=1}^k \frac{\nabla_{\epsilon} p_{[T^{-1}]}(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right] \right\} \\
&= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left\{ \left[\sum_{i=1}^k \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \right]^{-1} \left[\sum_{i=1}^k \frac{p(\Theta^i | \mathcal{D})}{q(\Theta^i)} \nabla_{\epsilon} \log p_{[T^{-1}]}(\Theta^i | \mathcal{D}) \right] \right\} \\
&= -\mathbb{E}_{\Theta^{1:k} \sim q(\Theta)} \left\{ \frac{1}{\tilde{\omega}} \sum_{i=1}^k \omega_i \left[\text{trace}(\mathcal{A}_p(\Theta^i; \mathcal{D})) \right] \right\}, \tag{26}
\end{aligned}$$

where $\omega_k = p(\Theta^i; \mathcal{D})/q(\Theta^i)$ and $\tilde{\omega} = \sum_{i=1}^k \omega_i$.

B Samples Updating for Stein VIWAE

let $\{\theta_j^{1:k,t}\}_{j=1}^M$ and $\{z_{jn}^{1:k,t}\}_{j=1}^M$ denote the samples acquired at iteration t of the learning procedure. To update samples of $\theta^{1:k}$, we apply the transformation $\theta_j^{(i,t+1)} = T(\theta_j^{(i,t)}; \mathcal{D}) = \theta_j^{(i,t)} + \epsilon \psi(\theta_j^{(i,t)}; \mathcal{D})$, for $i = 1, \dots, k$, by approximating the expectation by samples $\{z_{jn}^{1:k}\}_{j=1}^M$, and we have

$$\theta_j^{(i,t+1)} = \theta_j^{(i,t)} + \epsilon \Delta \theta_j^{(i,t)}, \text{ for } i = 1, \dots, k, \tag{27}$$

with

$$\Delta \theta_j^{(i,t)} \approx \frac{1}{M} \sum_{j'=1}^M \left[\frac{1}{\tilde{\omega}} \sum_{i'=1}^k \omega_i (\nabla_{\theta_{j'}^{(i',t)}} k_{\theta}(\theta_{j'}^{(i',t)}, \theta_j^{(i,t)})) + k(\theta_{j'}^{(i',t)}, \theta_j^{(i,t)}) \nabla_{\theta_{j'}^{(i',t)}} \log \tilde{p}(\theta_{j'}^{(i',t)}; \mathcal{D}) \right] \tag{28}$$

$$\omega_i \approx \frac{1}{M} \sum_{n=1}^N \sum_{j=1}^M \frac{p(\theta^i, z_{jn}^i, \mathbf{x}_n)}{q(\theta^i)q(z_{jn}^i)}, \quad \tilde{\omega} = \sum_{i=1}^k \omega_i \tag{29}$$

$$\nabla_{\theta} \log \tilde{p}(\theta; \mathcal{D}) \approx \frac{1}{M} \sum_{n=1}^N \sum_{j=1}^M \nabla_{\theta} \log p(\mathbf{x}_n | z_{jn}, \theta) p(\theta).$$

Similarly, when updating samples of the latent variables, we have

$$z_{jn}^{(i,t+1)} = z_{jn}^{(i,t)} + \epsilon \Delta z_{jn}^{(i,t)}, \text{ for } i = 1, \dots, k, \tag{30}$$

with

$$\Delta z_{jn}^{(i,t)} \approx \frac{1}{M} \sum_{j'=1}^M \left[\frac{1}{\tilde{\omega}_n} \sum_{i'=1}^k \omega_{in} (\nabla_{z_{j'n}^{(i',t)}} k_z(z_{j'n}^{(i',t)}, z_{jn}^{(i,t)})) + k_z(z_{j'n}^{(i',t)}, z_{jn}^{(i,t)}) \nabla_{z_{j'n}^{(i',t)}} \log \tilde{p}(z_{j'n}^{(i',t)}; \mathcal{D}) \right] \tag{31}$$

$$\omega_{in} \approx \frac{1}{M} \sum_{j=1}^M \frac{p(\theta^i, z_{jn}^i, \mathbf{x}_n)}{q(\theta^i)q(z_{jn}^i)}, \quad \tilde{\omega}_n = \sum_{i=1}^k \omega_{in} \tag{32}$$

$$\nabla_{z_n} \log \tilde{p}(z_n; \mathcal{D}) \approx \frac{1}{M} \sum_{j=1}^M \nabla_{z_n} \log p(\mathbf{x}_n | z_n, \theta'_j) p(z_n) \tag{33}$$

C Samples Updating for Semi-supervised Learning

The expectations in (13) and (14) in the main paper are approximated by samples. For updating samples of θ , we have

$$\theta_j^{(t+1)} = \theta_j^{(t)} + \epsilon_1 \Delta \theta_j^{(t)}, \tag{34}$$

with

$$\Delta\boldsymbol{\theta}_j^{(t)} \approx \frac{1}{M} \sum_{j'=1}^M [k_{\boldsymbol{\theta}}(\boldsymbol{\theta}_{j'}^{(t)}, \boldsymbol{\theta}_j^{(t)}) \nabla_{\boldsymbol{\theta}_{j'}} \log \tilde{p}(\boldsymbol{\theta}_{j'}^{(t)}; \mathcal{D}, \mathcal{D}_l) e + \nabla_{\boldsymbol{\theta}_{j'}} k_{\boldsymbol{\theta}}(\boldsymbol{\theta}_{j'}^{(t)}, \boldsymbol{\theta}_j^{(t)})] \quad (35)$$

$$\nabla_{\boldsymbol{\theta}} \log \tilde{p}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{D}_l) \approx \frac{1}{M} \sum_{j=1}^M \left\{ \sum_{\mathbf{x}_n \in \mathcal{D}} \nabla_{\boldsymbol{\theta}} \log p(\mathbf{x}_n | \mathbf{z}_{jn}, \boldsymbol{\theta}) + \sum_{\mathbf{x}_n \in \mathcal{D}_l} \nabla_{\boldsymbol{\theta}} \log p(\mathbf{x}_n | \mathbf{z}_{jn}, \boldsymbol{\theta}) \right\} p(\boldsymbol{\theta}). \quad (36)$$

Similarly, when updating samples of $\tilde{\boldsymbol{\theta}}$, we have

$$\tilde{\boldsymbol{\theta}}_j^{(t+1)} = \tilde{\boldsymbol{\theta}}_j^{(t)} + \epsilon_2 \Delta \tilde{\boldsymbol{\theta}}_j^{(t)}, \quad (37)$$

with

$$\begin{aligned} \Delta \tilde{\boldsymbol{\theta}}_j^{(t)} &\approx \frac{1}{M} \sum_{j'=1}^M [k_{\tilde{\boldsymbol{\theta}}}(\tilde{\boldsymbol{\theta}}_{j'}^{(t)}, \tilde{\boldsymbol{\theta}}_j^{(t)}) \nabla_{\tilde{\boldsymbol{\theta}}_{j'}} \log \tilde{p}(\tilde{\boldsymbol{\theta}}_{j'}^{(t)}; \mathcal{D}_l) + \nabla_{\tilde{\boldsymbol{\theta}}_{j'}} k_{\tilde{\boldsymbol{\theta}}}(\tilde{\boldsymbol{\theta}}_{j'}^{(t)}, \tilde{\boldsymbol{\theta}}_j^{(t)})] \\ \nabla_{\tilde{\boldsymbol{\theta}}} \log \tilde{p}(\tilde{\boldsymbol{\theta}}; \mathcal{D}_l) &\approx \frac{1}{M} \sum_{j=1}^M \sum_{\mathbf{y}_n \in \mathcal{D}_l} \nabla_{\tilde{\boldsymbol{\theta}}} \log p(\mathbf{y}_n | \mathbf{z}_{jn}, \tilde{\boldsymbol{\theta}}) p(\tilde{\boldsymbol{\theta}}). \end{aligned} \quad (38)$$

Similarly, samples of $\mathbf{z}_n \in \mathcal{Z}_l$ are updated

$$\mathbf{z}_{jn}^{(t+1)} = \mathbf{z}_{jn}^{(t)} + \epsilon \Delta \mathbf{z}_{jn}^{(t)}, \quad (39)$$

with

$$\Delta \mathbf{z}_{jn}^{(t)} = \frac{1}{M} \sum_{j'=1}^M [k_{\mathbf{z}}(\mathbf{z}_{j'n}^{(t)}, \mathbf{z}_{jn}^{(t)}) \nabla_{\mathbf{z}_{j'n}} \log \tilde{p}(\mathbf{z}_{j'n}^{(t)}; \mathcal{D}_l) + \nabla_{\mathbf{z}_{j'n}} k_{\mathbf{z}}(\mathbf{z}_{j'n}^{(t)}, \mathbf{z}_{jn}^{(t)})] \quad (40)$$

$$\nabla_{\mathbf{z}_n} \log \tilde{p}(\mathbf{z}_n; \mathcal{D}_l) \approx \frac{1}{M} \sum_{j=1}^M \nabla_{\mathbf{z}_n} p(\mathbf{z}_n) \left\{ \log p(\mathbf{x}_n | \mathbf{z}_n, \boldsymbol{\theta}_j') + \zeta \log p(\mathbf{y}_n | \mathbf{z}_n, \tilde{\boldsymbol{\theta}}_j') \right\}, \quad (41)$$

where ζ is a tuning parameter that balances the two components. Motivated by assigning the same weight to every data point [3], we set $\zeta = N_X / (C\rho)$ in the experiments, where N_X is the dimension of \mathbf{x}_n , C is the number of categories for the corresponding label and ρ is the proportion of labeled data in the mini-batch.

D Posterior of Gaussian Mixture Model

Consider $\mathbf{z} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \mathbf{I}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_2, \mathbf{I})$ and $\mathbf{x}_n \sim \mathcal{N}(\boldsymbol{\theta}\mathbf{z}, \sigma^2\mathbf{I})$, where $\mathbf{z} \in \mathbb{R}^K$, $\mathbf{x} \in \mathbb{R}^P$ and $\boldsymbol{\theta} \in \mathbb{R}^{P \times K}$. We have

$$\begin{aligned} p(\mathbf{z}|\mathbf{x}) &\propto p(\mathbf{x})p(\mathbf{z}) \propto \exp \left\{ -\frac{(\mathbf{x} - \boldsymbol{\theta}\mathbf{z})^T(\mathbf{x} - \boldsymbol{\theta}\mathbf{z})}{2\sigma^2} \right\} \\ &\quad \times \left\{ \exp \left\{ -\frac{(\mathbf{z} - \boldsymbol{\mu}_1)^T(\mathbf{z} - \boldsymbol{\mu}_1)}{2} \right\} + \exp \left\{ -\frac{(\mathbf{z} - \boldsymbol{\mu}_2)^T(\mathbf{z} - \boldsymbol{\mu}_2)}{2} \right\} \right\} \end{aligned} \quad (42)$$

$$\begin{aligned} &= \exp \left\{ -\frac{1}{2} \left[\mathbf{z}^T \left(\frac{\boldsymbol{\theta}^T \boldsymbol{\theta}}{\sigma^2} + \mathbf{I} \right) \mathbf{z} - 2 \left(\frac{\mathbf{y}^T \boldsymbol{\theta}}{\sigma^2} + \boldsymbol{\mu}_1 \right) \mathbf{z} + \frac{\mathbf{x}^T \mathbf{x}}{\sigma^2} + \boldsymbol{\mu}_1^T \boldsymbol{\mu}_1 \right] \right\} \\ &\quad + \exp \left\{ -\frac{1}{2} \left[\mathbf{z}^T \left(\frac{\boldsymbol{\theta}^T \boldsymbol{\theta}}{\sigma^2} + \mathbf{I} \right) \mathbf{z} - 2 \left(\frac{\mathbf{y}^T \boldsymbol{\theta}}{\sigma^2} + \boldsymbol{\mu}_2 \right) \mathbf{z} + \frac{\mathbf{x}^T \mathbf{x}}{\sigma^2} + \boldsymbol{\mu}_2^T \boldsymbol{\mu}_2 \right] \right\}. \end{aligned} \quad (43)$$

Let

$$\boldsymbol{\Sigma} = \frac{\boldsymbol{\theta}^T \boldsymbol{\theta}}{\sigma^2} + \mathbf{I}, \quad \hat{\boldsymbol{\mu}}_1 = \boldsymbol{\Sigma}^{-1} \left(\frac{\mathbf{y}^T \boldsymbol{\theta}}{\sigma^2} - \boldsymbol{\mu}_1 \right), \quad p_1 = \frac{\mathbf{x}^T \mathbf{x}}{\sigma^2} + \boldsymbol{\mu}_1^T \boldsymbol{\mu}_1 - \hat{\boldsymbol{\mu}}_1^T \boldsymbol{\Sigma} \hat{\boldsymbol{\mu}}_1, \quad (44)$$

$$\hat{\boldsymbol{\mu}}_2 = \boldsymbol{\Sigma}^{-1} \left(\frac{\mathbf{y}^T \boldsymbol{\theta}}{\sigma^2} - \boldsymbol{\mu}_2 \right), \quad p_2 = \frac{\mathbf{x}^T \mathbf{x}}{\sigma^2} + \boldsymbol{\mu}_2^T \boldsymbol{\mu}_2 - \hat{\boldsymbol{\mu}}_2^T \boldsymbol{\Sigma} \hat{\boldsymbol{\mu}}_2, \quad (45)$$

The density in (43) can be rewritten as

$$p(\mathbf{z}|\mathbf{x}) \propto \exp\{p_1\} \exp\left\{-\frac{1}{2}(\mathbf{z} - \hat{\boldsymbol{\mu}}_1)^T \boldsymbol{\Sigma}(\mathbf{z} - \hat{\boldsymbol{\mu}}_1)\right\} + \exp\{p_2\} \exp\left\{-\frac{1}{2}(\mathbf{z} - \hat{\boldsymbol{\mu}}_2)^T \boldsymbol{\Sigma}(\mathbf{z} - \hat{\boldsymbol{\mu}}_2)\right\}. \quad (46)$$

Therefore, we have $\mathbf{z}|\mathbf{x} \sim p(\mathbf{z}|\mathbf{x}) = \hat{p}\mathcal{N}(\hat{\boldsymbol{\mu}}_1, \boldsymbol{\Sigma}) + (1 - \hat{p})\mathcal{N}(\hat{\boldsymbol{\mu}}_2, \boldsymbol{\Sigma})$, where

$$\hat{p} = \frac{1}{1 + \exp(p_2 - p_1)}. \quad (47)$$

E Model Architecture

Table 1: Architecture of the models for semi-supervised classification on ImageNet. BN denotes batch normalization. The layer in bracket indicates the number of layers stacked.

Output Size	Encoder	Decoder
224 × 224 × 4 for encoder 224 × 224 × 3 for decoder	RGB image x_n stacked by ξ	RGB image x_n
56 × 56 × 64	7×7 conv, 64 kernels, LeakyRelu, stride 4, BN $\left[3 \times 3 \text{ conv, 64 kernels, LeakyRelu, stride 1, BN} \right] \times 3$	
28 × 28 × 128	3×3 conv, 128 kernels, LeakyRelu, stride 2, BN $\left[3 \times 3 \text{ conv, 128 kernels, LeakyRelu, stride 1, BN} \right] \times 3$	
14 × 14 × 256	3×3 conv, 256 kernels, LeakyRelu, stride 2, BN $\left[3 \times 3 \text{ conv, 256 kernels, LeakyRelu, stride 1, BN} \right] \times 3$	
7 × 7 × 512	3×3 conv, 512 kernels, LeakyRelu, stride 2, BN $\left[3 \times 3 \text{ conv, 512 kernels, LeakyRelu, stride 1, BN} \right] \times 3$	
	latent code z_n	
	1 × 1 conv, 2048 kernels, LeakyRelu	
	average pooling, 1000-dimentional fully connected layer	
	softmax, label y_n	

F Additional Results

Gaussian Mixture Model Figure 1 and 2 show the performance of Stein VAE approximations for the true posterior using $M = 10$, $M = 20$, $M = 50$ and $M = 100$ samples on test data.

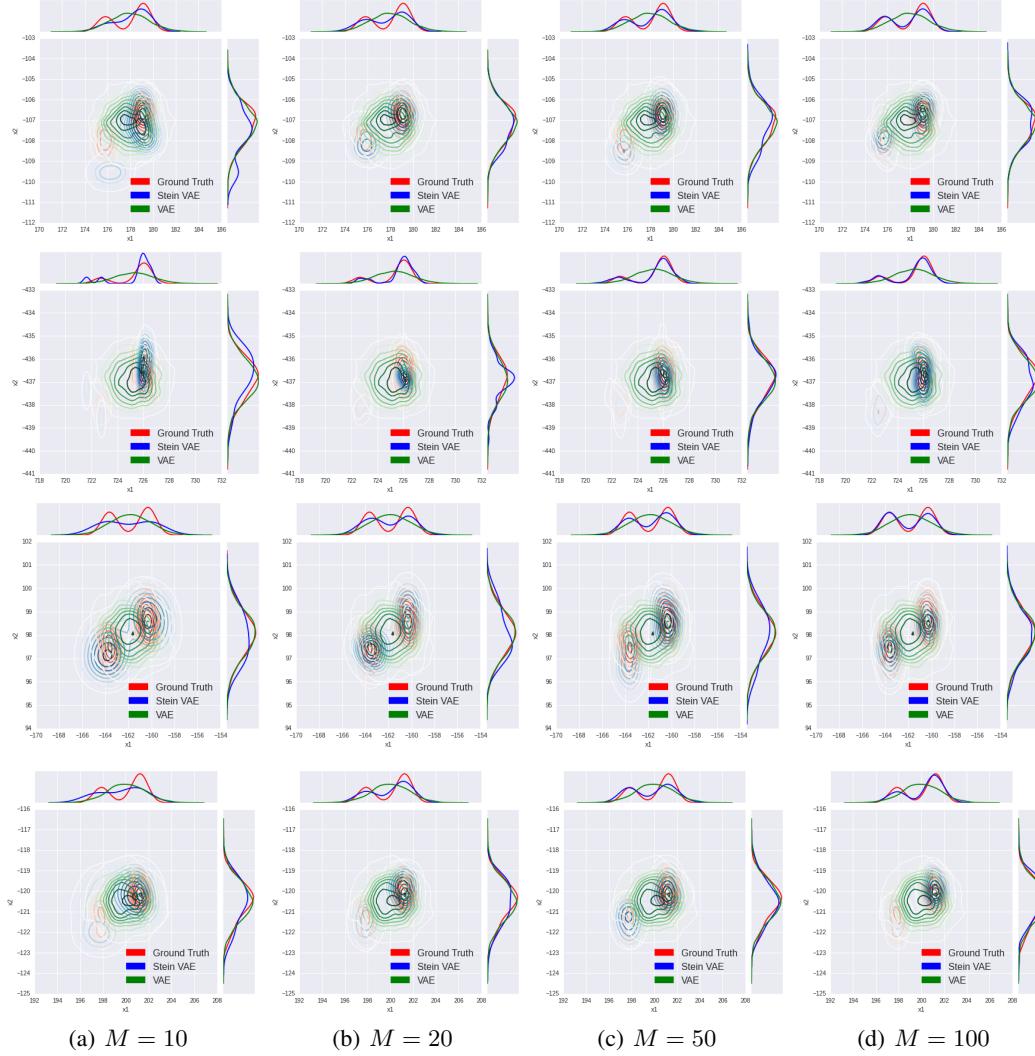


Figure 1: Approximation of posterior distribution: Stein VAE vs. VAE. The figures represent different samples of Stein VAE. Each row corresponds to the same test data, and each column corresponds to the same number of samples with (a) 10 samples; (b) 20 samples; (c) 50 samples; (d) 100 samples.

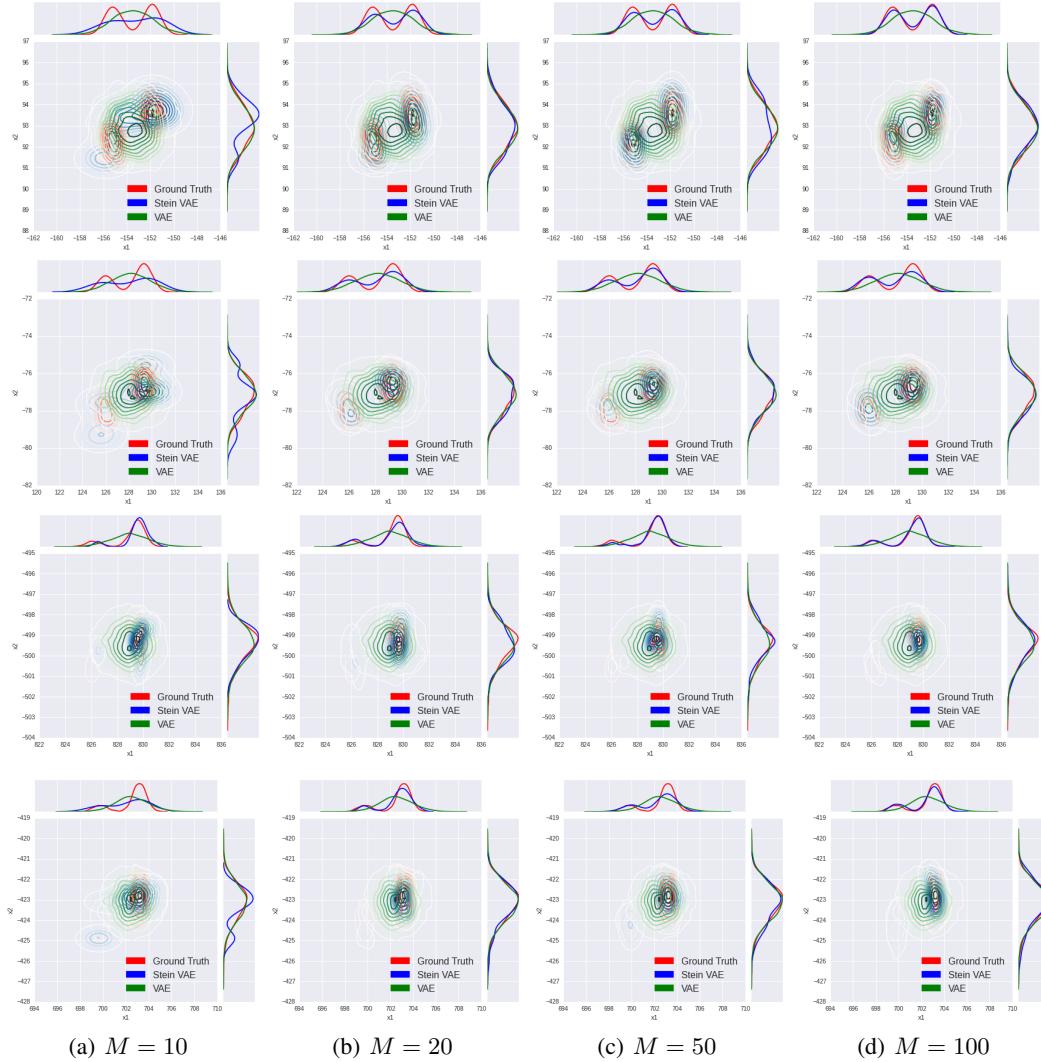


Figure 2: Approximation of posterior distribution: Stein VAE vs. VAE. The figures represent different samples of Stein VAE. Each row corresponds to the same test data, and each column corresponds to the same number of samples with (a) 10 samples; (b) 20 samples; (c) 50 samples; (d) 100 samples.

Poisson Factor Analysis We show the marginal and pairwise posteriors of test data in Figure 3.

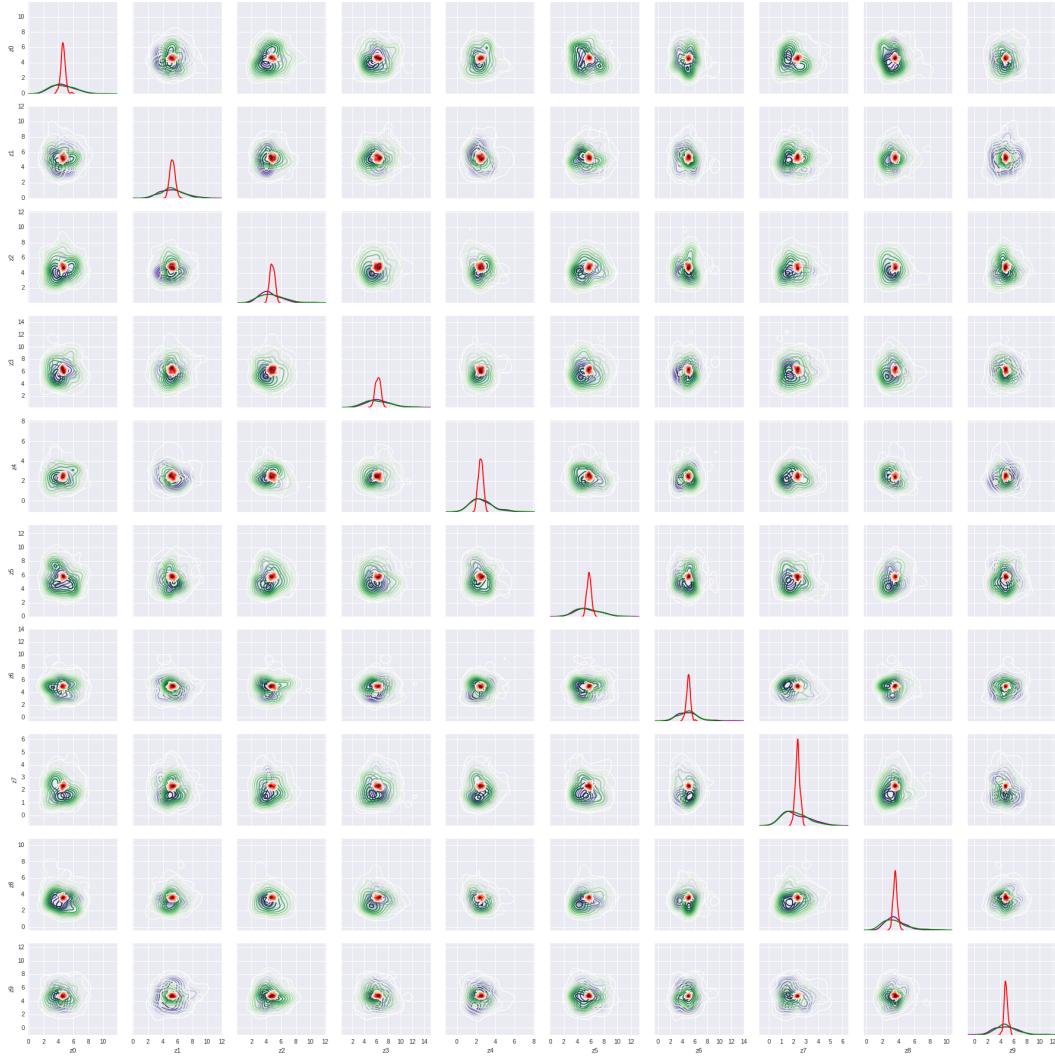


Figure 3: Univariate marginals and pairwise posteriors. Purple, red and green represent the distribution inferred from MCMC, standard VAE and Stein VAE, respectively.

References

- [1] Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. In *ICLR*, 2016.
- [2] Q. Liu and D. Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *NIPS*, 2016.
- [3] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, A. Stevens, and L. Carin. Variational autoencoder for deep learning of images, labels and captions. In *NIPS*, 2016.