Statistical Guarantees for Variational Autoencoders using PAC-Bayesian Theory

Sokhna Diarra Mbacke, Florence Clerc, Pascal Germain

NeurIPS 2023 Spotlight

Summary

- We derive the first PAC-Bayes bound for conditional posterior distributions.
- We use this result to derive statistical guarantees for VAEs.
- Our results include the reconstruction, regeneration, and generation guarantees for the standard VAE.

PAC-Bayesian Theory

- PAC-Bayes is powerful tool in statistical learning theory.
- PAC-Bayes has been applied to a multitude of problems.
- Our first result is a novel PAC-Bayes bound with a conditional posterior distribution.

Next up

1. General PAC-Bayes Bound with a Conditional Posterior
2. Variational Autoencoders
3. Reconstruction Guarantees
4. Regeneration and Generation Guarantees
5. Conclusion

The General theorem: Definitions

- $(\mathcal{X}, d)$ is a metric space;
- $\mu \in \mathcal{M}_+^1(\mathcal{X})$ is the data-generating distribution;
- $S = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \overset{\text{iid}}{\sim} \mu$ is a set of observed samples;
- $\mathcal{H}$ is the hypothesis class;
- $p(h) \in \mathcal{M}_+^1(\mathcal{H})$ is the prior distribution on $\mathcal{H}$;
- $\lambda > 0$ and $\delta \in (0, 1)$;

General PAC-Bayes Bound with a Conditional Posterior

- The goal is to obtain a PAC-Bayes bound for conditional posterior distributions $q(h|\mathbf{x})$, conditioned on elements of the instance space.
- The main goal for this bound is the analysis of VAEs, since the variational posterior $q_\phi(\mathbf{z}|\mathbf{x}_i)$ is conditional.
- This result requires the following assumption.

Assumption 1

Assumption

We say that a distribution $q(h|\mathbf{x})$ and a loss function ℓ satisfy Assumption 1 with a constant $K > 0$ if there exists a family $\mathcal{E}$ of functions $\mathcal{H} \to \mathbb{R}$ such that the following properties hold.

1. The function $\mathbf{x} \mapsto q(\cdot|\mathbf{x})$ is continuous in the following sense: for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$,
$$d_{\mathcal{E}}\left(q(h|\mathbf{x}_1), q(h|\mathbf{x}_2)\right) \leq K d(\mathbf{x}_1, \mathbf{x}_2).$$
2. For any $\mathbf{x} \in \mathcal{X}$, the function $\ell(\cdot, \mathbf{x}) : \mathcal{H} \to \mathbb{R}$ is in $\mathcal{E}$:
$$\ell(\cdot, \mathbf{x}) \in \mathcal{E}, \quad \text{for any } \mathbf{x} \in \mathcal{X}.$$

The General PAC-Bayes Bound for Conditional Posteriors

Theorem

With probability $1-\delta$ *over the random draw of* $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \overset{iid}{\sim} \mu$, *the following holds for any conditional posterior* $q(h|\mathbf{x})$ *satisfying Assumption 1:*

$$\mathop{\mathbb{E}}_{\mathbf{x}\sim\mu} \mathop{\mathbb{E}}_{h\sim q(h|\mathbf{x})} \ell(h,\mathbf{x}) \leq \frac{1}{n}\sum_{i=1}^{n} \mathop{\mathbb{E}}_{h\sim q(h|\mathbf{x}_i)} \ell(h,\mathbf{x}_i) + \frac{1}{\lambda}\sum_{i=1}^{n} \mathrm{KL}(q(h|\mathbf{x}_i)\,||\,p(h)) +$$

$$\frac{K}{n}\sum_{i=1}^{n} \mathop{\mathbb{E}}_{\mathbf{x}\sim\mu} d(\mathbf{x},\mathbf{x}_i) + \frac{1}{\lambda}\log\frac{1}{\delta} +$$

$$\frac{n}{\lambda}\log \mathop{\mathbb{E}}_{\mathbf{x}\sim\mu} \mathop{\mathbb{E}}_{h\sim p(h)} e^{\frac{\lambda}{n}\left(\mathbb{E}_{\mathbf{x}'\sim\mu}[\ell(h,\mathbf{x}')]-\ell(h,\mathbf{x})\right)}$$

Next up

1. General PAC-Bayes Bound with a Conditional Posterior
2. Variational Autoencoders
3. Reconstruction Guarantees
4. Regeneration and Generation Guarantees
5. Conclusion

Variational Autoencoders

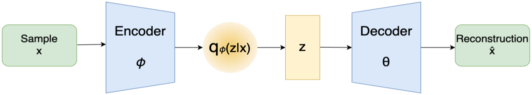

- **The latent space.** $\mathcal{Z} = \mathbb{R}^{d_{\mathcal{Z}}}$
- **The encoder.** $Q_\phi : \mathcal{X} \to \mathbb{R}^{2d_{\mathcal{Z}}}$, where $Q_\phi(\mathbf{x}) = \begin{bmatrix} \mu_\phi(\mathbf{x}) \\ \sigma_\phi(\mathbf{x}) \end{bmatrix}$.
- **The variational posterior.** $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}\left(\mu_\phi(\mathbf{x}), \mathrm{diag}(\sigma^2_\phi(\mathbf{x}))\right)$.
- **The decoder.** $g_\theta : \mathcal{Z} \to \mathcal{X}$.
- **The reconstruction loss.** $\ell^\theta_{\mathrm{rec}}(\mathbf{z}, \mathbf{x}) = \|\mathbf{x} - g_\theta(\mathbf{z})\|$.
- **Lipschitz norms.** $\|Q_\phi\|_{\mathrm{Lip}} = K_\phi$ and $\|g_\theta\|_{\mathrm{Lip}} = K_\theta$.

Variational Autoencoders: Optimization Objective

Given a training set $S = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, the encoder and decoder networks are jointly trained by minimizing the following objective:

$$\mathcal{L}_{\mathsf{VAE}}(\phi, \theta) = \frac{1}{n} \sum_{i=1}^{n} \left[\underbrace{\mathop{\mathbb{E}}_{\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x}_i)} \ell^{\theta}_{\mathsf{rec}}(\mathbf{z}, \mathbf{x}_i)}_{\text{Reconstruction loss}} + \beta \underbrace{\mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z}))}_{\text{KL loss}} \right].$$

Applying the General Theorem to VAEs: Assumption 1

Recall Assumption 1: There exists $\mathcal{E} \subseteq \mathbb{R}^{\mathcal{Z}}$ such that:

1. For any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$, $\quad d_{\mathcal{E}}\left(q(\mathbf{z}|\mathbf{x}_1), q(\mathbf{z}|\mathbf{x}_2)\right) \leq K d(\mathbf{x}_1, \mathbf{x}_2)$;
2. For any $\mathbf{x} \in \mathcal{X}$, $\quad \ell(\cdot, \mathbf{x}) \in \mathcal{E}$.

Applying the General Theorem to VAEs: Assumption 1

Recall Assumption 1: There exists $\mathcal{E} \subseteq \mathbb{R}^{\mathcal{Z}}$ such that:

1. For any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$, $\quad d_{\mathcal{E}}\left(q(\mathbf{z}|\mathbf{x}_1), q(\mathbf{z}|\mathbf{x}_2)\right) \leq K d(\mathbf{x}_1, \mathbf{x}_2)$;
2. For any $\mathbf{x} \in \mathcal{X}$, $\quad \ell(\cdot, \mathbf{x}) \in \mathcal{E}$.

Proposition

Consider a VAE with parameters ϕ and θ and let $K_\phi, K_\theta \in \mathbb{R}$ be the Lipschitz norms of the encoder and decoder respectively. Then the variational distribution $q_\phi(\mathbf{z}|\mathbf{x})$ satisfies Assumption 1, with $\mathcal{E} = \mathrm{Lip}_{K_\theta}(\mathcal{Z}, \mathbb{R})$, $\ell = \ell^\theta_{rec}$, and $K = K_\phi K_\theta$.

Next up

1. General PAC-Bayes Bound with a Conditional Posterior
2. Variational Autoencoders
3. Reconstruction Guarantees
4. Regeneration and Generation Guarantees
5. Conclusion

Reconstruction Guarantees for Bounded Instance Spaces

Theorem

Assuming $\Delta = \sup_{\mathbf{x},\mathbf{x}' \in \mathcal{X}} d(\mathbf{x}, \mathbf{x}') < \infty$, *with probability at least* $1 - \delta$, *the following holds:*

$$\mathop{\mathbb{E}}_{\mathbf{x}\sim\mu} \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x})} \ell^\theta_{rec}(\mathbf{z}, \mathbf{x}) \leq \frac{1}{n} \sum_{i=1}^{n} \left\{ \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x}_i)} \ell^\theta_{rec}(\mathbf{z}, \mathbf{x}_i) \right\} + \frac{1}{\lambda} \sum_{i=1}^{n} \mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z})) + \frac{1}{\lambda} \log \frac{1}{\delta} + K_\phi K_\theta \Delta + \frac{\lambda \Delta^2}{8n}.$$

Reconstruction Guarantees Under the Manifold Assumption

Theorem

Assuming $\mu = g^* \sharp p^*$, *where* $g^* \in \mathrm{Lip}_{K_*}$ *and* $p^* = \mathcal{N}(\mathbf{0}, \mathbf{I})$ *on* $\mathbb{R}^{d^*}$, *with probability at least* $1 - \delta - \frac{nd^*}{2} e^{-a^2/2}$, *the following holds for any posterior* $q_\phi(\mathbf{z}|\mathbf{x})$*:*

$$\mathop{\mathbb{E}}_{\mathbf{x}\sim\mu} \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x})} \ell^\theta_{rec}(\mathbf{z}, \mathbf{x}) \leq \frac{1}{n} \sum_{i=1}^{n} \left\{ \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x}_i)} \ell^\theta_{rec}(\mathbf{z}, \mathbf{x}_i) \right\} + \frac{1}{\lambda} \sum_{i=1}^{n} \mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z}))$$
$$+ \frac{1}{\lambda} \log \frac{1}{\delta} + K_\phi K_\theta K_* \sqrt{(1 + a^2)d^*} + \frac{\lambda K_*^2}{2n}.$$

Next up

1 General PAC-Bayes Bound with a Conditional Posterior

2 Variational Autoencoders

3 Reconstruction Guarantees

4 Regeneration and Generation Guarantees

5 Conclusion

The VAE's generative model

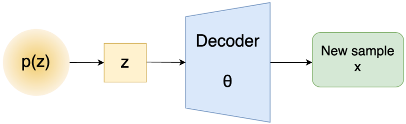

- Once trained, the VAE defines a generative model: $g_\theta \sharp p(\mathbf{z})$.
- Our goal is to bound the distance: $W_1(\mu, g_\theta \sharp p(\mathbf{z}))$.
- Considering the regenerated distribution: $\hat{\mu}_{\phi,\theta} = \frac{1}{n}\sum_{i=1}^{n} g_\theta \sharp q_\phi(\mathbf{z}|\mathbf{x}_i)$, we use the inequality:

$$W_1(\mu, g_\theta \sharp p(\mathbf{x})) \leq W_1(\mu, \hat{\mu}_{\phi,\theta}) + W_1(\hat{\mu}_{\phi,\theta}, g_\theta \sharp p(\mathbf{x})).$$

Regeneration Guarantees for Bounded Instance Spaces

Theorem

With probability at least $1-\delta$*, the following holds:*

$$W_1(\mu, \hat{\mu}_{\phi,\theta}) \leq \frac{1}{n}\sum_{i=1}^{n}\left\{ \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x}_i)} \ell^{\theta}_{rec}(\mathbf{z}, \mathbf{x}_i)\right\} + \frac{1}{\lambda}\sum_{i=1}^{n}\mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i)\,||\,p(\mathbf{z})) + \frac{1}{\lambda}\log\frac{1}{\delta} + \frac{\lambda\Delta^2}{8n}.$$

Generation Guarantees for Bounded Instance Spaces

Theorem

With probability at least $1-\delta$, *the following holds:*

$$W_1(\mu, g_\theta \sharp p(\mathbf{z})) \leq \frac{1}{n}\sum_{i=1}^{n}\left\{\underset{q_\phi(\mathbf{z}|\mathbf{x}_i)}{\mathbb{E}} \ell_{rec}^{\theta}(\mathbf{z}, \mathbf{x}_i)\right\} + \frac{1}{\lambda}\sum_{i=1}^{n}\mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z}))$$
$$+\frac{1}{\lambda}\log\frac{1}{\delta} + \frac{\lambda\Delta^2}{8n} + \frac{K_\theta}{n}\sum_{i=1}^{n}\sqrt{\left\|\mu_\phi(\mathbf{x}_i)\right\|^2 + \left\|\sigma_\phi(\mathbf{x}_i) - \vec{1}\right\|^2}.$$

Regeneration Guarantees under the Manifold Assumption

Theorem

With probability at least $1 - \delta$, the following holds:

$$W_1(\mu, \hat{\mu}_{\phi,\theta}) \leq \frac{1}{n} \sum_{i=1}^{n} \left\{ \mathop{\mathbb{E}}_{q_\phi(\mathbf{z}|\mathbf{x}_i)} \ell_{rec}^{\theta}(\mathbf{z}, \mathbf{x}_i) \right\} + \frac{1}{\lambda} \sum_{i=1}^{n} \mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z})) + \frac{1}{\lambda} \log \frac{1}{\delta} + \frac{\lambda K_*^2}{2n}.$$

Generation Guarantees under the Manifold Assumption

Theorem

With probability at least $1-\delta$, the following holds:

$$W_1(\mu, g_\theta \sharp p(\mathbf{z})) \leq \frac{1}{n}\sum_{i=1}^{n}\left\{\underset{q_\phi(\mathbf{z}|\mathbf{x}_i)}{\mathbb{E}} \ell_{rec}^{\theta}(\mathbf{z}, \mathbf{x}_i)\right\} + \frac{1}{\lambda}\sum_{i=1}^{n}\mathrm{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) \,||\, p(\mathbf{z}))$$
$$+\frac{1}{\lambda}\log\frac{1}{\delta} + \frac{\lambda K_*^2}{2n} + \frac{K_\theta}{n}\sum_{i=1}^{n}\sqrt{\|\mu_\phi(\mathbf{x}_i)\|^2 + \left\|\sigma_\phi(\mathbf{x}_i) - \vec{1}\right\|^2}.$$

Next up

1. General PAC-Bayes Bound with a Conditional Posterior
2. Variational Autoencoders
3. Reconstruction Guarantees
4. Regeneration and Generation Guarantees
5. Conclusion

Conclusion

- We proved, to the best of our knowledge, the first statistical guarantees for VAEs.
- We consider the standard VAE, with no additional noise on the parameters.
- Our results cover the reconstruction, regeneration and generation properties of VAEs.
- The seamless integration of VAE and PAC-Bayes is promising.