
Supplementary Material of Quantized Variational Inference

A ELBO derivation

Assumes that we have observations y , latent variables z and a model $p(y, z)$ with p the density function for the distribution y . By Bayes' Theorem

$$\begin{aligned} p(z|y) &= \frac{p(y|z)p(z)}{p(y)} \\ &= \frac{p(y|z)p(z)}{\int_z p(z, y) dz}. \end{aligned}$$

Using the definition of KL divergence,

$$\begin{aligned} \text{KL}[q_\lambda(z) \| p(z|y)] &= \int_z q_\lambda(z) \log \frac{q_\lambda(z)}{p(z|y)} dz \\ &= - \int_z q_\lambda(z) \log \frac{p(z|y)}{q_\lambda(z)} dz \\ &= - \int_z q_\lambda(z) \log \frac{p(z, y)}{q_\lambda(z)} dz + \int_z q_\lambda(z) \log p(y) dz \\ &= - \int_z q_\lambda(z) \log \frac{p(z, y)}{q_\lambda(z)} dz + \log p(y) \int_z q_\lambda(z) dz \\ &= -\mathcal{L}(\lambda) + \log p(y). \end{aligned}$$

Rearranging the terms gives equation (1).

B Proofs

Let $f(X) \in L^2_{\mathbb{R}^d}(\Omega, \mathcal{A}, \mathbb{P})$ and $X^{\Gamma_N, \lambda}$ the the optimal quantizer of X^λ . The general framework of our study can be stated as estimating the quantity

$$I = \mathbb{E}[f(X)]. \tag{1}$$

We define the *MC* and *OQ* estimators as

$$I_{MC} = \frac{1}{N} \sum_{i=1}^N f(X_i), \tag{2}$$

$$I_{OQ} = \sum_{i=1}^N \underbrace{\mathbb{P}(X^{\Gamma_N, \lambda} = x_i)}_{\omega_i} f(x_i). \tag{3}$$

It is direct to derive $\|I - I_{MC}\|_2 = \mathcal{O}(N^{-\frac{1}{2}})$. In the following we establish the approximation error for the I_{OQ} estimator.

In this part we demonstrates proposition 1 and proposition 2. The former is particularly important since it establishes an asymptomatic bound on the error produced by using QVI. When considering

it along with proposition 1 justifies QVI, for ranking models with it will produce true ranking provided that the relative difference in ELBO is lower than the quantization error. In the following we formally demonstrate such result (thorough investigation of optimal quantizer can be found in [8, 7]). We begin with the definition of a stationnary quantizer.

Definition 1. Let $\Gamma_N = \{x_1, \dots, x_N\}$ be a quantization scheme of X^λ . $X^{\Gamma_N, \lambda}$ is said to be stationary quantizer if the Voronoi partition induced by Γ_N satisfies $\mathbb{P}(X \in C_i(x)) > 0 \forall i \in \{1, \dots, N\}$ and

$$\mathbb{E} \left[X^\lambda | X^{\Gamma_N, \lambda} \right] = X^{\Gamma_N, \lambda}.$$

One of the first question raised by using optimal quantization $\mathbb{E} [H(X^{\Gamma_N, \lambda})]$ in place for $\mathbb{E} [H(X^\lambda)]$ is the error produced by such substitution. Let us remind that we denote $\widehat{\mathcal{L}}_{OQ}^N(\lambda) = \mathbb{E} [H(X^{\Gamma_N, \lambda})]$ the quantized ELBO estimator and $\mathcal{L}(\lambda) = \mathbb{E} [H(X^\lambda)]$ the true ELBO.

Lemma 1. Let $X^\lambda \in L^2_{\mathbb{R}^d}(\Omega, \mathcal{A}, \mathbb{P})$ and a H a continuous lipschitz function with Lipschitz constant C , we have

$$\left| \mathcal{L}(\lambda) - \widehat{\mathcal{L}}_{OQ}^N(\lambda) \right| \leq C \left\| X^\lambda - X^{\Gamma_N, \lambda} \right\|_2.$$

Proof.

$$\left| \mathbb{E} [H(X^\lambda)] - \mathbb{E} [H(X^{\Gamma_N, \lambda})] \right| \leq \mathbb{E} \left[\left| H(X^\lambda) - H(X^{\Gamma_N, \lambda}) \right| | X^{\Gamma_N, \lambda} \right] \quad (4)$$

$$\begin{aligned} &\leq C \left\| X^\lambda - X^{\Gamma_N, \lambda} \right\|_1 \\ &\leq C \left\| X^\lambda - X^{\Gamma_N, \lambda} \right\|_2. \end{aligned} \quad (5)$$

We use Jensen inequality in equation 4 and the monoticity of the $L_p(\Omega, \mathcal{A}, \mathbb{P})$ norm as a function of p in equation 5. \square

Proposition 1. Let $X^\lambda \in L^2_{\mathbb{R}^d}(\Omega, \mathcal{A}, \mathbb{P})$ and $X^{\Gamma_N, \lambda}$ the associated optimal quantizer, under the hypothesis that H is a convex lipschitz function,

$$\widehat{\mathcal{L}}_{OQ}^N(\lambda) \leq \mathcal{L}(\lambda).$$

Proof.

$$\begin{aligned} \widehat{\mathcal{L}}_{OQ}^N(\lambda) &= \mathbb{E} [H(X^{\Gamma_N, \lambda})] \\ &= \mathbb{E} \left[H \left(\mathbb{E} [X^\lambda | X^{\Gamma_N, \lambda}] \right) \right] \end{aligned} \quad (6)$$

$$\begin{aligned} &\leq \mathbb{E} \left[\mathbb{E} [H(X^\lambda) | X^{\Gamma_N, \lambda}] \right] \\ &= \mathbb{E} [H(X^\lambda)] \\ &= \mathcal{L}(\lambda) \end{aligned} \quad (7)$$

When we used Lemma 1 in equation 6 and the conditional Jensen inequality to obtain 7. \square

Proposition 2. Let $\lambda^* = \min_{\lambda \in \mathbb{R}^K} \mathcal{L}(\lambda)$ and $\lambda_q^* = \min_{\lambda \in \mathbb{R}^K} \widehat{\mathcal{L}}_{OQ}^N(\lambda)$. Under the same assumptions than proposition 1,

$$\mathcal{L}(\lambda^*) - \widehat{\mathcal{L}}_{OQ}^N(\lambda_q^*) \leq C \left[2 \left\| X^{\lambda^*} - X^{\Gamma, \lambda^*} \right\|_2 + \left\| X^{\lambda_q^*} - X^{\Gamma, \lambda_q^*} \right\|_2 \right].$$

Proof. A immediate consequence of proposition 1 is that $\widehat{\mathcal{L}}_{OQ}^N(\lambda_q^*) \leq \mathcal{L}(\lambda^*)$. Then, we can write

$$\begin{aligned}\mathcal{L}(\lambda^*) - \widehat{\mathcal{L}}_{OQ}^N(\lambda_q^*) &= \mathcal{L}(\lambda^*) - \widehat{\mathcal{L}}_{OQ}^N(\lambda^*) \\ &\quad + \widehat{\mathcal{L}}_{OQ}^N(\lambda^*) - \mathcal{L}(\lambda_q^*) \\ &\quad + \mathcal{L}(\lambda_q^*) - \widehat{\mathcal{L}}_{OQ}^N(\lambda_q^*) \\ &\leq C \|X^{\lambda^*} - X^{\Gamma, \lambda^*}\|_2 \\ &\quad + C \|X^{\lambda_q^*} - X^{\Gamma, \lambda_q^*}\|_2 \\ &\quad + C \|X^{\lambda^*} - X^{\Gamma, \lambda^*}\|_2\end{aligned}$$

Using Lemma 1 and noting that

$$\widehat{\mathcal{L}}_{OQ}^N(\lambda^*) - \mathcal{L}(\lambda_q^*) \leq \widehat{\mathcal{L}}_{OQ}^N(\lambda^*) - \mathcal{L}(\lambda^*),$$

proposition 2 follows. \square

Finally, Zador's theorem is used to derive non-asymptotic bound (see [5] for a complete proof).

Theorem 1 (Zador's Theorem). *Let $X^\lambda \in L^2_{\mathbb{R}^d}(\Omega, \mathcal{A}, \mathbb{P})$ and $X^{\Gamma, \lambda}$ the associated optimal quantizer at level N , there exists a real constant $C_{d,p}$ such that*

$$\forall N \geq 1, \quad \|X - \widehat{X}^{\Gamma, \lambda}\|_p \leq C_{d,p} N^{-\frac{1}{d}}$$

Where $C_{d,p}$ depends only d and p . This result can be vastly improved when H exhibits more regularity. For instance, if H is an α hölderian function, we can obtain a bound in $\mathcal{O}(N^{-\frac{1+\alpha}{d}})$ [7].

C Experiments

Bayesian Linear Regression. We used three different real-world dataset, namely Forests Fire, Boston housing datasets from the UCI repository [2] and Life Expectancy dataset from the Global Health Observatory repository. The generative Bayesian Linear Gaussian Model used is as follow.

$$\begin{aligned}\mathbf{b}_i &\sim \mathcal{N}(\mu_\beta, \sigma_\beta), & \text{intercepts} \\ y_i &\sim \mathcal{N}(\mathbf{x}_i^\top \mathbf{b}_i, \epsilon), & \text{output}\end{aligned}$$

Let D be the dimension of the feature space. The dimension of the parameter space for a gaussian variational distribution under the mean-field assumption is $K = 2D$.

Poisson Generalized Linear Model. The frisk dataset is a record of stops and searches practice on civilians in New York City for fifteen months in 1998 – 1999. It contains information about locations, ethnicity and crime statistics for each area. The question is whether these stops targeted particular groups after taking into account population and crime rates in each group for a particular precinct.

We can trace back the use of Poisson Generalized Linear Model for this use case to [3]. The model writes as follow

$$\mu \sim \mathcal{N}(0, 10^2) \quad \text{mean offset} \quad (8)$$

$$\log \sigma_\alpha^2, \log \sigma_\beta^2 \sim \mathcal{N}(0, 10^2) \quad \text{group variances} \quad (9)$$

$$\alpha_e \sim \mathcal{N}(0, \sigma_\alpha^2) \quad \text{ethnicity effect} \quad (10)$$

$$\beta_p \sim \mathcal{N}(0, \sigma_\beta^2) \quad \text{precinct effect} \quad (11)$$

$$\log \lambda_{ep} = \mu + \alpha_e + \beta_p + \log N_{ep} \quad \text{log rate} \quad (12)$$

$$Y_{ep} \sim \text{Poisson}(\lambda_{ep}) \quad \text{stops events} \quad (13)$$

$$(14)$$

Y_{ep} denotes the number of frisk events for the ethnic group e in the precinct p . N_{ep} is the number of arrests for the ethnic group e in the precinct p . Hence, in this model, α_e and α_p represents the ethnicity and precinct effect. The dataset contains three ethnicities and thirty-two precinct, which therefore exhibits $K = 70$ variational parameters.

Bayesian Neural Network. The Bayesian Neural Network (BNN) consists of a Multi Layer Perceptron (MLP) ψ of 30 ReLU activated neurons with normal prior weights and inverse Gamma hyperprior on the mean and variance. Regression is performed on the metro dataset.

$$\alpha \sim \text{Gamma}(1, 0.1) \quad \text{weights hyper prior} \quad (15)$$

$$\tau \sim \text{Gamma}(1, 0.1) \quad \text{group variances} \quad (16)$$

$$w \sim \mathcal{N}\left(0, \frac{1}{\alpha}\right), \quad \text{neural network weights} \quad (17)$$

$$y \sim \mathcal{N}\left(\psi(w, x), \frac{1}{\tau}\right) \quad \text{output} \quad (18)$$

D Thanks to open source libraries

This work and many others would have been impossible without free, open-source computational frameworks and libraries. We particularly acknowledge Python 3 [9], Tensorflow [1], Numpy [6] and Matplotlib [4].

References

- [1] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*. 2015.
- [2] Dheeru Dua and Casey Graff. *UCI Machine Learning Repository*. University of California, Irvine, School of Information and Computer Sciences. 2017.
- [3] Andrew Gelman, Jeffrey Fagan, and Alex Kiss. “An Analysis of the New York City Police Department’s “Stop-and-Frisk” Policy in the Context of Claims of Racial Bias”. en. In: *Journal of the American Statistical Association* 102.479 (Sept. 2007), pp. 813–823.
- [4] John D Hunter. “Matplotlib: A 2D Graphics Environment”. In: *Computing in science & engineering* 9.3 (2007), pp. 90–95.
- [5] Harald Luschgy and Gilles Pagès. “Functional Quantization Rate and Mean Regularity of Processes with an Application to Lévy Processes”. EN. In: *Annals of Applied Probability* 18.2 (Apr. 2008), pp. 427–469.
- [6] Travis E Oliphant. *A Guide to NumPy*. Vol. 1. Trelgol Publishing USA, 2006.
- [7] Gilles Pagès. “Introduction to Vector Quantization and Its Applications for Numerics”. en. In: *ESAIM: Proceedings and Surveys* 48 (Jan. 2015), pp. 29–79.
- [8] Gilles Pagès. *Numerical Probability: An Introduction with Applications to Finance*. en. Universitext. Springer International Publishing, 2018.
- [9] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009.