The Thermodynamic Variational Objective

Vaden Masrani, Tuan Anh Le, Frank Wood
vadmas@cs.ubc.ca, tuananh@mit.edu, fwood@cs.ubc.ca

THE UNIVERSITY OF BRITISH COLUMBIA
Massachusetts Institute of Technology

TL;DR

- **The TVO is a new objective for training both continuous and discrete deep generative models that is as broadly applicable as the ELBO.**
- **The TVO achieves state-of-the-art model and inference network learning *without* using the reparameterization trick.**
- **The TVO is a generalization of the objectives used in variational inference [5], variational autoencoders [7,8], wake sleep [3], and inference compilation [6]**
- **The TVO arises from a novel connection between *thermodynamic integration* (TI) and *variational inference (VI).***

Thermodynamic Integration and Variational Inference

TI refresher

- TI is a technique from physics to calculate the log ratio of two unknown normalizing constants [1, 2].
- Consider two densities with intractable normalizing constants

$$\pi_1(\mathbf{z}) = \frac{\tilde{\pi}_1(\mathbf{z})}{Z_1} \qquad \pi_0(\mathbf{z}) = \frac{\tilde{\pi}_0(\mathbf{z})}{Z_0}$$

- Form a geometric path between $\tilde{\pi}_1(\mathbf{z})$ and $\tilde{\pi}_0(\mathbf{z})$ using scalar parameter β

$$\pi_\beta(\mathbf{z}) := \frac{\tilde{\pi}_\beta(\mathbf{z})}{Z_\beta} = \frac{\tilde{\pi}_1(\mathbf{z})^\beta \tilde{\pi}_0(\mathbf{z})^{1-\beta}}{Z_\beta}$$

- Then compute $\log(Z_1/Z_0)$ using the central TI identity:

$$\log(Z_1) - \log(Z_0) = \int_0^1 \mathbb{E}_{\pi_\beta}\left[\frac{\partial}{\partial\beta}\log\tilde{\pi}_\beta(\mathbf{z})\right] d\beta$$

- The integrand is typically evaluated at multiple β points using MCMC, then approximated using numerical methods
- $\log\tilde{\pi}_\beta(\mathbf{z})$ is referred to as the "potential" in physics

Connecting TI + VI

- Consider the model and inference network

$$\tilde{\pi}_1(\mathbf{z}) := p_\theta(\mathbf{x}, \mathbf{z}) \qquad Z_1 = p_\theta(\mathbf{x})$$
$$\tilde{\pi}_0(\mathbf{z}) := q_\phi(\mathbf{z} \mid \mathbf{x}) \qquad Z_0 = 1$$

- Form a geometric path between $p_\theta(\mathrm{x}, \mathrm{z})$ and $q_\phi(\mathrm{z} \mid \mathrm{x})$ using scalar parameter β

$$\pi_\beta(\mathbf{z}) := \frac{\tilde{\pi}_\beta(\mathbf{z})}{Z_\beta} = \frac{p_\theta(\mathbf{x}, \mathbf{z})^\beta q_\phi(\mathbf{z} \mid \mathbf{x})^{1-\beta}}{Z_\beta}$$

- With this path, the first derivative of the potential is the "instantaneous ELBO"!

$$\frac{\partial}{\partial\beta}\log\tilde{\pi}_\beta(\mathbf{z}) = \log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}$$

- Plugging in to ◆, we arrive at the *thermodynamic variational identity* (TVI)

$$\log p_\theta(\mathbf{x}) - 0 = \int_0^1 \mathbb{E}_{\pi_\beta}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right] d\beta$$

Is the ELBO at $\beta = 0$

Is the EUBO at $\beta = 1$

Integrand strictly increasing

The Thermodynamic Variational Objective (TVO)

The TVO is a *K-term* Riemann integral approximation to $\log p_\theta(\mathbf{x})$. The ELBO is a *1-term* Riemann integral approximation to $\log p_\theta(\mathbf{x})$.

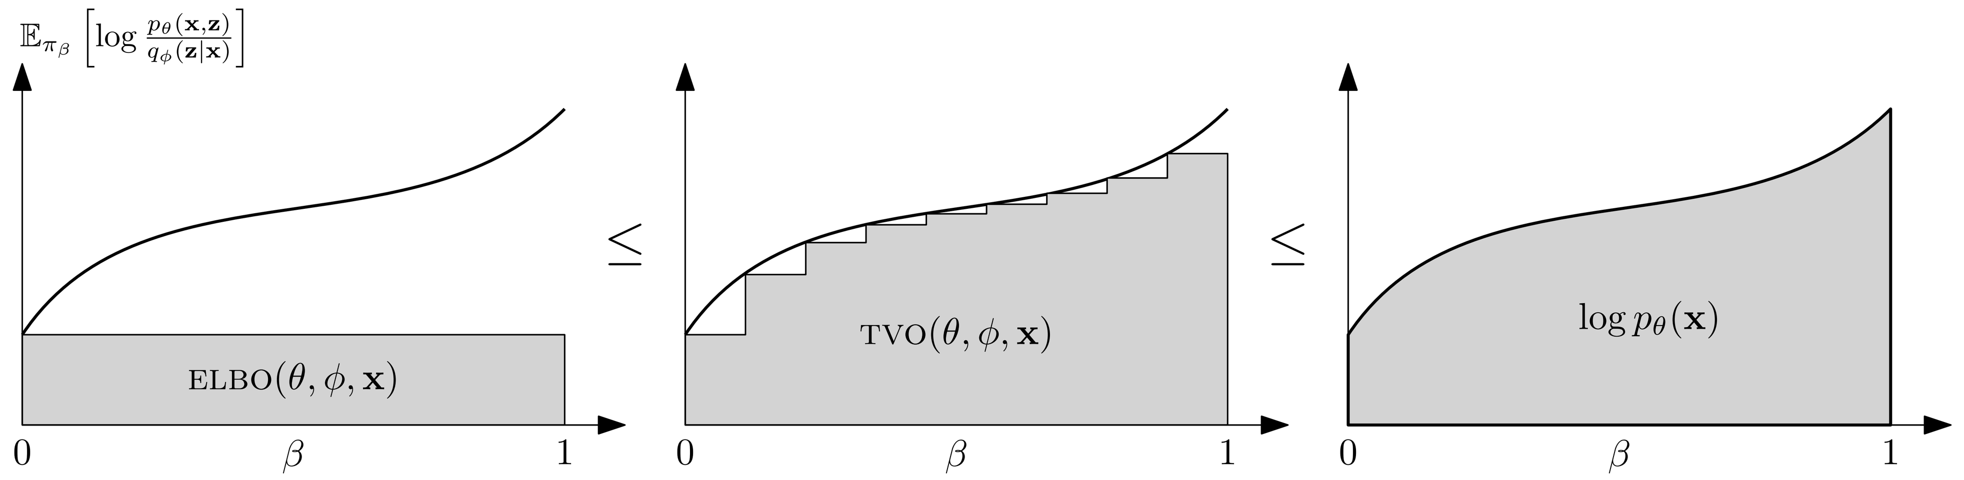

$$\underbrace{\frac{1}{K}\left[\mathrm{ELBO}(\theta, \phi, \mathbf{x}) + \sum_{k=1}^{K-1}\mathbb{E}_{\pi_{\beta_k}}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right]\right]}_{\mathrm{TVO}(\theta, \phi, \mathbf{x})} \leq \underbrace{\int_0^1 \mathbb{E}_{\pi_\beta}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right] d\beta = \log p_\theta(\mathbf{x})}_{\text{The Thermodynamic Variational Identity (TVI)}}$$

Optimizing the TVO

- To use the TVO as an objective function, we need to be able to compute expectations and gradients
- **Expectations** under $\pi_\beta(\mathbf{z})$ can be computed **simply by exponentiating importance weights by β**

$$\frac{\tilde{\pi}_\beta(\mathbf{x}, \mathbf{z}_s)}{q_\phi(\mathbf{z}_s \mid \mathbf{x})} = \frac{p_\theta(\mathbf{x}, \mathbf{z}_s)^\beta q_\phi(\mathbf{z}_s \mid \mathbf{x})^{1-\beta}}{q_\phi(\mathbf{z}_s \mid \mathbf{x})} = \left(\frac{p_\theta(\mathbf{x}, \mathbf{z}_s)}{q_\phi(\mathbf{z}_s \mid \mathbf{x})}\right)^\beta = w_s^\beta$$

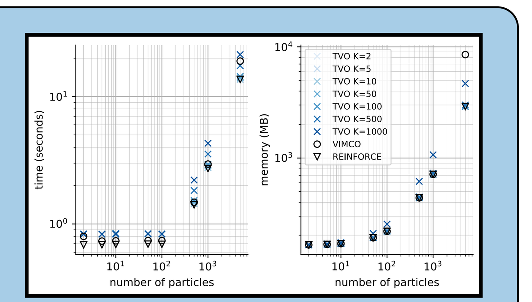

Partitions are cheap (because we can reuse samples!)

- **Gradients** can be computed using the score-function based ***covariance gradient estimator***

Unnormalized, easy to compute
(unlike reinforce: $\nabla_\theta \mathbb{E}_{\pi_\theta}[f(\mathbf{z})] = \mathbb{E}_{\pi_\theta}[f(\mathbf{z})\nabla_\theta \log \pi_\theta(\mathbf{z})]$)

$$\nabla_\theta \mathbb{E}_{\pi_\theta}[f(\mathbf{z})] = \mathrm{Cov}_{\pi_\theta}\left[f(\mathbf{z}), \nabla_\theta \log\tilde{\pi}_\theta(\mathbf{z})\right]$$
$$= \mathbb{E}_{\pi_\theta}\left[\left(f(\mathbf{z}) - \mathbb{E}_{\pi_\theta}\left[f(\mathbf{z})\right]\right)\left(\nabla_\theta \log\tilde{\pi}_\theta(\mathbf{z}) - \mathbb{E}_{\pi_\theta}\left[\nabla_\theta \log\tilde{\pi}_\theta(\mathbf{z})\right]\right)\right]$$

Equivalent to "average" baseline commonly used w/ reinforce

Replacement for $\nabla_\theta \log Z_\theta$ using similar derivation as ◆

Experiments

Discrete

Model: sigmoid belief network
Dataset: MNIST
Baselines: VIMCO [9], reweighted wake sleep (RWS) [4]

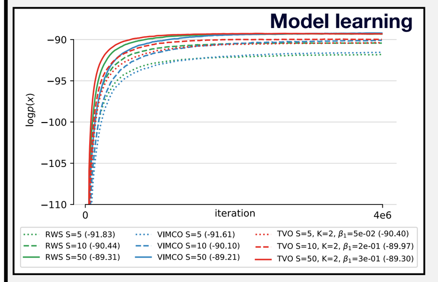

TVO (red) achieves SOTA in terms of convergence rate and final test log evidence for S=5, 10. At S = 50 VIMCO achieves a higher test log evidence but converges more slowly.

Inference networks learned with the TVO are better - i.e. closer to $p_\theta(\mathbf{z} \mid \mathbf{x})$ as measured by the KL divergence - than VIMCO and RWS.

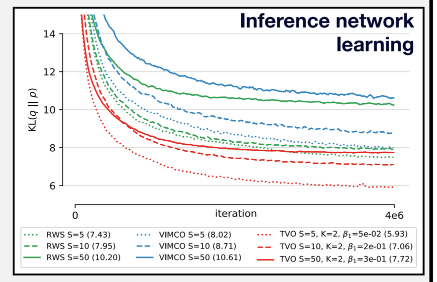

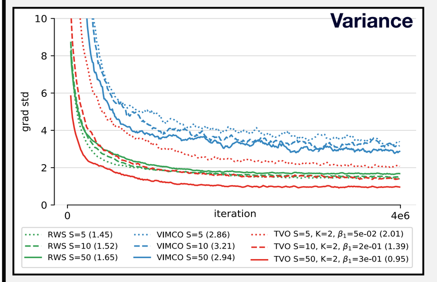

Standard deviation of the gradient estimator for each objective. TVO has lowest variance, VIMCO is highest variance, RWS is in the middle.

Continuous

Model: Variational autoencoder
Dataset: MNIST
Baselines: Variational Autoencoder [7,8], Importance-weighted Auto Encoder [4]

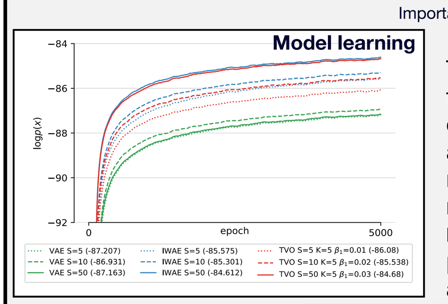

The TVO outperforms the VAE and performs competitively with IWAE at 50 samples, despite not using the reparameterization trick. IWAE is the top performing objective in all cases.

Standard deviation of the gradient estimator for each objective. The TVO has lower variance than IWAE but higher than VAE.

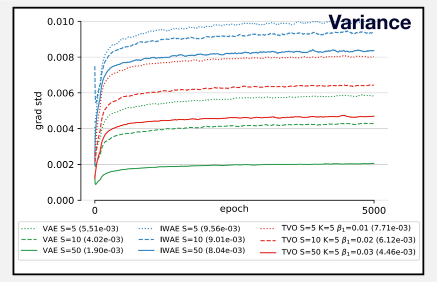

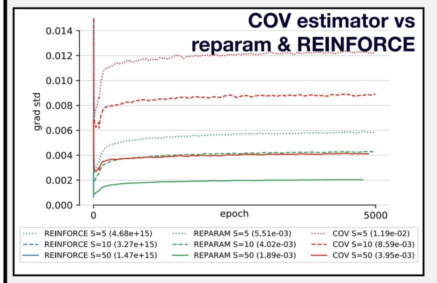

The covariance estimator has higher variance than the reparameterization trick for all S but much lower than REINFORCE, which is numerically unstable. (ELBO objective used in all cases)

We also investigated the effect of $\beta_{1:K}$ locations

We have empirically observed the integrand has a point of maximum curvature, β^*

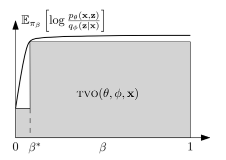

$\beta_1 \ll \beta^*$

$\beta_1 \approx \beta^*$

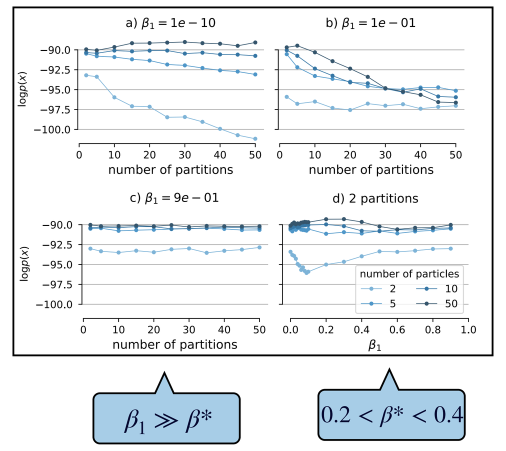

$\beta_1 \gg \beta^*$

$0.2 < \beta^* < 0.4$

Generalizing Existing Objective

- The left and right endpoints of the TVI correspond to the ELBO and EUBO, respectively

$$\mathrm{ELBO} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right] \qquad \mathrm{EUBO} = \mathbb{E}_{p_\theta(\mathbf{z}|\mathbf{x})}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right]$$

- Using a right Riemann sum results in an upper bound to $\log p_\theta(\mathbf{x})$ that can be minimized

$$\mathrm{TVO}^U(\theta, \phi, \mathbf{x}) := \frac{1}{K}\left[\mathrm{EUBO}(\theta, \phi, \mathbf{x}) + \sum_{k=1}^{K-1}\mathbb{E}_{\pi_{\beta_k}}\left[\log\frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{z} \mid \mathbf{x})}\right]\right] \geq \log p_\theta(\mathbf{x})$$

- Methods that optimize the ELBO or EUBO are 1-term left / right Riemann approximations to the TVI

ELBO objectives

- Variational Inference [5]
- Variational AutoEncoders [7,8]
- Wake in WS [3]

EUBO objectives

- Inference Compilation [6]
- Sleep in WS [3]
- Sleep-ϕ in RWS [4]

References

[1] Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, pages 163–185, 1998.
[2] Radford M Neal. Probabilistic inference using Markov chain Monte Carlo methods. 1993.
[3] Geoffrey E Hinton, Peter Dayan, Brendan J Frey, and Radford M Neal. The "wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158–1161, 1995
[4] Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. In International Conference on Learning Representations, 2015.
[5] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. Journal of the American Statistical Association, 112(518):859–877, 2017
[6] Tuan Anh Le, Atilim Gunes Baydin, and Frank Wood. Inference compilation and universal probabilistic programming. arXiv preprint arXiv:1610.09900, 2016.
[7] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In International Conference on Learning Representations, 2014.
[8] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In International Conference on Machine Learning, 2014.
[9] Andriy Mnih and Danilo Rezende. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, pages 2188–2196, 2016.
[10] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In *International Conference on Learning Representations*, 2016.