
Variational Laplace Autoencoders

Yookoon Park¹ Chris Dongjoo Kim¹ Gunhee Kim¹

Abstract

Variational autoencoders (Kingma & Welling, 2014) employ an amortized inference model to approximate the posterior of latent variables. However, such amortized variational inference faces two challenges: (1) the limited posterior expressiveness of fully-factorized Gaussian assumption and (2) the amortization error of the inference model. We present a novel approach that addresses both challenges. First, we focus on ReLU networks with Gaussian output and illustrate their connection to probabilistic PCA. Building on this observation, we derive an iterative algorithm that finds the mode of the posterior and apply full-covariance Gaussian posterior approximation centered on the mode. Subsequently, we present a general framework named *Variational Laplace Autoencoders* (VLAEs) for training deep generative models. Based on the *Laplace approximation* of the latent variable posterior, VLAEs enhance the expressiveness of the posterior while reducing the amortization error. Empirical results on MNIST, Omniglot, Fashion-MNIST, SVHN and CIFAR10 show that the proposed approach significantly outperforms other recent amortized or iterative methods on the ReLU networks.

1. Introduction

Variational autoencoders (VAEs) (Kingma & Welling, 2014) are deep latent generative models which have been popularly used in various domains of data such as images, natural language and sound (Gulrajani et al., 2017; Gregor et al., 2015; Bowman et al., 2016; Chung et al., 2015; Roberts et al., 2018). VAEs introduce an amortized inference network (*i.e.* an *encoder*) to approximate the posterior distribution of latent variable \mathbf{z} and maximize the evidence lower-bound (ELBO) of the data. However, the two major limitations of

VAEs are: (1) the constrained expressiveness of the fully-factorized Gaussian posterior assumption and (2) the amortization error (Cremer et al., 2018) of the inference model due to the dynamic posterior prediction.

There have been various attempts to address these problems. A representative line of works belongs to the category of normalizing flows (Rezende & Mohamed, 2015; Kingma et al., 2016; Tomczak & Welling, 2016), which apply a chain of invertible transformation with tractable densities in order to represent a more flexible posterior distribution. However, not only do they incur additional parameter overhead for the inference model but are yet prone to the amortization error as they entirely depend on the dynamic inference.

Recently, iterative approaches based on gradient-based refinement of the posterior parameters have been proposed (Krishnan et al., 2018; Kim et al., 2018; Marino et al., 2018). These methods aim to reduce the amortization error by augmenting the dynamic inference with an additional inner-loop optimization of the posterior. Nonetheless, they still rely on the fully-factorized Gaussian assumption, and accordingly fail to enhance the expressiveness of the posterior.

We develop a novel approach that tackles both challenges by (1) iteratively updating the mode of the approximate posterior and (2) defining a full-covariance Gaussian posterior centered on the mode, whose covariance is determined by the local behavior of the generative network. By deducing the approximate posterior directly from the generative network, not only are we able to minimize the amortization error but also model the rich correlations between the latent variable dimensions.

We start from the class of neural networks of rectified linear activations (*e.g.* ReLU) (Montufar et al., 2014; Pascanu et al., 2014) and Gaussian output, which is universally popular for modeling continuous data such images (Krizhevsky et al., 2012; Gregor et al., 2015) and also as a building block of deep latent models (Rezende et al., 2014; Sønderby et al., 2016). Subsequently, we present a generalized framework named *Variational Laplace Autoencoders* (VLAEs), which encompasses the general class of differentiable neural networks. We show that the ReLU networks are in fact a special case of VLAEs that admits efficient computation. In addition, we illustrate an example for Bernoulli output networks.

¹Neural Processing Research Center, Seoul National University, Seoul, South Korea. Correspondence to: Gunhee Kim <gunhee@snu.ac.kr>.

In summary, the contributions of this work are as follows:

- We relate Gaussian output ReLU networks to probabilistic PCA (Tipping & Bishop, 1999) and thereby derive an iterative update for finding the mode of the posterior and a full-covariance Gaussian posterior approximation at the mode.
- We present *Variational Laplace Autoencoders* (VLAEs), a general framework for training deep generative models based on the *Laplace approximation* of the latent variable posterior. VLAEs not only minimize the amortization error but also provide the expressive power of full-covariance Gaussian, with no additional parameter overhead for the inference model. To the best of our knowledge, this work is the first attempt to apply the Laplace approximation to the training of deep generative models.
- We evaluate our approach on five benchmark datasets of MNIST, Omniglot, Fashion-MNIST, SVHN and CIFAR-10. Empirical results show that VLAEs bring significant improvement over other recent amortized or iterative approaches, including VAE (Kingma & Welling, 2014), semi-amortized VAE (Kim et al., 2018; Marino et al., 2018; Krishnan et al., 2018) and VAE with Householder Flow (Tomczak & Welling, 2016).

2. Background

2.1. Variational Autoencoders

For the latent variable model $p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\theta}(\mathbf{x}|\mathbf{z})p(\mathbf{z})$, variational inference (VI) (Hinton & Camp, 1993; Waterhouse et al., 1996; Jordan et al., 1999) fits an approximate distribution $q(\mathbf{z}; \boldsymbol{\lambda})$ to the intractable posterior $p_{\theta}(\mathbf{z}|\mathbf{x})$ and maximizes the *evidence lower-bound* (ELBO):

$$\mathcal{L}_{\theta}(\mathbf{x}; \boldsymbol{\lambda}) = \mathbb{E}_{q(\mathbf{z}; \boldsymbol{\lambda})}[\ln p_{\theta}(\mathbf{x}, \mathbf{z}) - \ln q(\mathbf{z}; \boldsymbol{\lambda})] \quad (1)$$

$$= \ln p_{\theta}(\mathbf{x}) - D_{KL}(q(\mathbf{z}; \boldsymbol{\lambda})||p_{\theta}(\mathbf{z}|\mathbf{x})) \quad (2)$$

$$\leq \ln p_{\theta}(\mathbf{x}), \quad (3)$$

where $q(\mathbf{z}; \boldsymbol{\lambda})$ is assumed to be a simple distribution such as diagonal Gaussian and $\boldsymbol{\lambda}$ is the variational parameter of the distribution (*e.g.* the mean and covariance). The learning involves first finding the optimal variational parameter $\boldsymbol{\lambda}^*$ that minimizes the variational gap $D_{KL}(q(\mathbf{z}; \boldsymbol{\lambda})||p_{\theta}(\mathbf{z}|\mathbf{x}))$ between the ELBO and the true marginal log-likelihood, and then updating the generative model θ using $\mathcal{L}_{\theta}(\mathbf{x}; \boldsymbol{\lambda}^*)$.

Variational autoencoders (VAE) (Kingma & Welling, 2014) amortize the optimization problem of $\boldsymbol{\lambda}$ using the inference model ϕ that dynamically predicts the approximate posterior as a function of \mathbf{x} :

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\ln p_{\theta}(\mathbf{x}, \mathbf{z}) - \ln q_{\phi}(\mathbf{z}|\mathbf{x})]. \quad (4)$$

The generative model (*decoder*) and the inference model (*encoder*) are jointly optimized. While such amortized variational inference (AVI) is highly efficient, there remain two fundamental challenges: (1) the limited expressiveness of the approximate posterior and (2) the amortization error, both of which will be discussed in the following sections.

2.2. Limited Posterior Expressiveness

VAEs approximate the posterior with the fully-factorized Gaussian $q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\phi}(\mathbf{x}), \text{diag}(\boldsymbol{\sigma}_{\phi}^2(\mathbf{x})))$. However the fully-factorized Gaussian may fail to accurately capture the complex true posterior distribution, causing the *approximation error* (Cremer et al., 2018). As the ELBO (Eq.(4)) tries to reduce the gap $D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z}|\mathbf{x}))$, it will force the true posterior $p_{\theta}(\mathbf{z}|\mathbf{x})$ to match the fully-factorized Gaussian $q_{\phi}(\mathbf{z}|\mathbf{x})$, which negatively affect the capacity of the generative model (Mescheder et al., 2017).

Hence, it is encouraged to use more expressive families of distributions; for example, one natural expansion is to model the full-covariance matrix $\boldsymbol{\Sigma}$ of the Gaussian. However, it requires $O(D^2)$ variational parameters to be predicted, placing a heavy burden on the inference model.

Normalizing flows (Rezende & Mohamed, 2015; Kingma et al., 2016; Tomczak & Welling, 2016) approach this issue by using a class of invertible transformations whose densities are relatively easy to compute. However, the flow-based methods have drawbacks in that they incur additional parameter overhead for the inference model and are prone to the amortization error described below.

2.3. The Amortization Error

Another problem of VAEs stems from the nature of amortized inference where the variational parameter $\boldsymbol{\lambda}$ is not explicitly optimized, but is dynamically predicted by the inference model. The error of the dynamic inference is referred to as the *amortization error* and is closely related to the performance of the generative model (Cremer et al., 2018). The suboptimal posterior predictions loosen the bound in the ELBO (Eq.(2)) and result in biased gradient signals flowing to the generative parameters.

Recently, Kim et al. (2018) and Marino et al. (2018) address this issue by iteratively updating the predicted variational parameters using gradient-based optimization. However, they still rely on the fully-factorized Gaussian assumption, limiting the expressive power of the posterior.

3. Approach

We first describe probabilistic PCA (section 3.1) and piecewise linear ReLU networks (section 3.2). Based on the local linearity of such networks, we derive an iterative approach

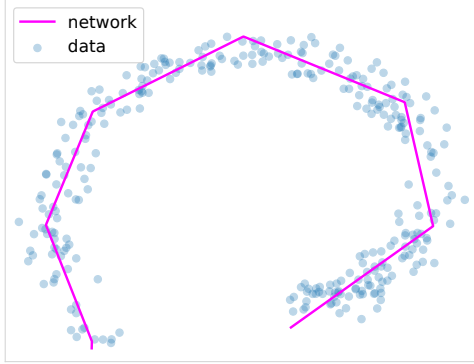


Figure 1. VAE’s manifold learned on 2D toy data using a 1D latent variable. The VAE with rectified linear activation learns a piecewise linear manifold and locally performs probabilistic PCA.

to find the mode of the posterior and define a full-covariance Gaussian posterior at the mode (section 3.3). Finally, we present a general framework for training deep generative models using the Laplace approximation of the latent variable posterior (section 3.4).

3.1. Probabilistic PCA

Probabilistic PCA (Tipping & Bishop, 1999) relates latent variable \mathbf{z} to data \mathbf{x} through linear mapping \mathbf{W} as

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (5)$$

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{W}\mathbf{z} + \mathbf{b}, \sigma^2\mathbf{I}), \quad (6)$$

where \mathbf{W} , \mathbf{b} , σ are the parameters to be learned. Note that this is basically a *linear* version of VAEs. Under this particular model, the posterior distribution of \mathbf{z} given \mathbf{x} can be computed in a closed form (Tipping & Bishop, 1999):

$$p_{\theta}(\mathbf{z}|\mathbf{x}) = \mathcal{N}\left(\frac{1}{\sigma^2}\mathbf{\Sigma}\mathbf{W}^T(\mathbf{x} - \mathbf{b}), \mathbf{\Sigma}\right), \quad (7)$$

$$\text{where } \mathbf{\Sigma} = \left(\frac{1}{\sigma^2}\mathbf{W}^T\mathbf{W} + \mathbf{I}\right)^{-1}. \quad (8)$$

3.2. Piece-wise Linear Neural Networks

Consider a following ReLU network $\mathbf{y} = g_{\theta}(\mathbf{z})$:

$$\mathbf{h}_{l+1} = \text{ReLU}(\mathbf{W}_l\mathbf{h}_l + \mathbf{b}_l), \text{ for } l = 0, \dots, L-1 \quad (9)$$

$$\mathbf{y} = \mathbf{W}_L\mathbf{h}_L + \mathbf{b}_L, \quad (10)$$

where $\mathbf{h}_0 = \mathbf{z}$, $\text{ReLU}(\mathbf{x}) = \max(\mathbf{0}, \mathbf{x})$ and L is the number of layers.

Our motivation is based on the observation that neural networks of rectified linear activations (*e.g.* ReLU, Leaky ReLU, Maxout) are *piece-wise linear* (Pascanu et al., 2014; Montufar et al., 2014). That is, the network segments the input space into linear regions within which it locally behaves

as a linear function:

$$g_{\theta}(\mathbf{z} + \boldsymbol{\epsilon}) \approx \mathbf{W}_{\mathbf{z}}(\mathbf{z} + \boldsymbol{\epsilon}) + \mathbf{b}_{\mathbf{z}}, \quad (11)$$

where the subscripts denote the dependence on \mathbf{z} .

To see this, note that applying a ReLU activation is equivalent to multiplying a corresponding mask matrix \mathbf{O} :

$$\text{ReLU}(\mathbf{W}\mathbf{x} + \mathbf{b}) = \mathbf{O}(\mathbf{W}\mathbf{x} + \mathbf{b}), \quad (12)$$

where \mathbf{O} is a diagonal matrix whose diagonal element o_i defines the activation pattern (Pascanu et al., 2014):

$$o_i = \begin{cases} 1 & \text{if } \mathbf{w}_i^T\mathbf{x} + b_i > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The set of input \mathbf{x} that satisfies the activation pattern $\{o_i\}_{i=1}^d$ such that $\{\mathbf{x} | \mathbb{I}(\mathbf{w}_i^T\mathbf{x} + b_i > 0) = o_i, \text{ for } i = 1, \dots, d\}$ defines a convex polytope as it is an intersection of half-spaces. Accordingly, within this convex polytope the mask matrix \mathbf{O} is constant.

We can obtain $\mathbf{W}_{\mathbf{z}}$ and $\mathbf{b}_{\mathbf{z}}$ in Eq.(11) by computing the activation masks during the forward pass and recursively multiplying them with the network weights:

$$\mathbf{y} = \mathbf{W}_L\text{ReLU}(\mathbf{W}_{L-1}\mathbf{h}_{L-1} + \mathbf{b}_{L-1}) + \mathbf{b}_L \quad (14)$$

$$= \mathbf{W}_L\mathbf{O}_{L-1}(\mathbf{W}_{L-1}\mathbf{h}_{L-1} + \mathbf{b}_{L-1}) + \mathbf{b}_L \quad (15)$$

$$= \mathbf{W}_L\mathbf{O}_{L-1}\mathbf{W}_{L-1}\cdots\mathbf{O}_0\mathbf{W}_0\mathbf{z} + \dots \quad (16)$$

$$= \mathbf{W}_{\mathbf{z}}\mathbf{z} + \mathbf{b}_{\mathbf{z}}. \quad (17)$$

Note that $\mathbf{W}_{\mathbf{z}}$ is the Jacobian of the network $\partial g_{\theta}(\mathbf{z})/\partial \mathbf{z}$. Similar results apply to other kinds of piece-wise linear activations such as Leaky ReLU (Maas et al., 2013) and MaxOut (Goodfellow et al., 2013).

3.3. Posterior Inference for Piece-wise Linear Networks

Consider the following nonlinear latent generative model:

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (18)$$

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(g_{\theta}(\mathbf{z}), \sigma^2\mathbf{I}), \quad (19)$$

where $g_{\theta}(\mathbf{z})$ is the ReLU network. In general, such nonlinear model does not allow the analytical computation of the posterior. Instead of using the amortized prediction $q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\phi}(\mathbf{x}), \text{diag}(\boldsymbol{\sigma}_{\phi}^2(\mathbf{x})))$ like VAEs, we present a novel approach that exploits the piece-wise linearity of generative networks.

The results in the previous section hints that the ReLU network $g_{\theta}(\mathbf{z})$ learns the piece-wise linear manifold of the data as illustrated in Fig.1, meaning that the model is locally equivalent to the probabilistic PCA. Based on this observation, we propose a new approach for posterior approximation which consists of two parts: (1) find the mode of

Algorithm 1 Posterior inference for piece-wise linear nets

Input: data \mathbf{x} , piece-wise linear generative network g_θ , inference network enc_ϕ , update steps T , decay α_t
Output: full-covariance Gaussian posterior $q(\mathbf{z}|\mathbf{x})$
 $\boldsymbol{\mu}_0 = \text{enc}_\phi(\mathbf{x})$
for $t = 0$ **to** $T - 1$ **do**
 Compute linear approximation $g_\theta(\mathbf{z}) \approx \mathbf{W}_t \mathbf{z} + \mathbf{b}_t$
 $\boldsymbol{\Sigma}_t \leftarrow (\sigma^{-2} \mathbf{W}_t^T \mathbf{W}_t + \mathbf{I})^{-1}$
 $\boldsymbol{\mu}' \leftarrow \sigma^{-2} \boldsymbol{\Sigma}_t \mathbf{W}_t^T (\mathbf{x} - \mathbf{b}_t)$
 $\boldsymbol{\mu}_{t+1} \leftarrow (1 - \alpha_t) \boldsymbol{\mu}_t + \alpha_t \boldsymbol{\mu}'$
end for
 Compute linear approximation $g_\theta(\mathbf{z}) \approx \mathbf{W}_T \mathbf{z} + \mathbf{b}_T$
 $\boldsymbol{\Sigma}_T \leftarrow (\sigma^{-2} \mathbf{W}_T^T \mathbf{W}_T + \mathbf{I})^{-1}$
 $q(\mathbf{z}|\mathbf{x}) \leftarrow \mathcal{N}(\boldsymbol{\mu}_T, \boldsymbol{\Sigma}_T)$
Return $q(\mathbf{z}|\mathbf{x})$

posterior where probability density is mostly concentrated, and (2) apply local linear approximation of the generative network (Eq.(11)) at the mode and analytically compute the posterior using the results of probabilistic PCA (Eq.(7)). Algorithm 1 outlines the proposed method.

We derive an update equation for the posterior mode exploiting the local linearity of ReLU networks. The results of probabilistic PCA (Eq.(7)) leads to the solution for the posterior mode under the linear model $\mathbf{y} = \mathbf{W}\mathbf{z} + \mathbf{b}$. Based on this insight, we first assume linear approximation to the generative network $g_\theta(\mathbf{z}) \approx \mathbf{W}_t \mathbf{z} + \mathbf{b}_t$ at the current estimate $\boldsymbol{\mu}_t$ at step t and update our mode estimate using the solution (Eq.(7)) under this linear model:

$$\boldsymbol{\mu}_{t+1} = \frac{1}{\sigma^2} \boldsymbol{\Sigma}_t \mathbf{W}_t^T (\mathbf{x} - \mathbf{b}_t), \quad (20)$$

where $\boldsymbol{\Sigma}_t$ is defined as in Eq.(8).

To take advantage of the efficiency of the amortized inference, we initialize the estimate using an inference model (encoder) as $\boldsymbol{\mu}_0 = \text{enc}_\phi(\mathbf{x})$ and iterate the update (Eq.(20)) for T steps. Fig. 2 illustrates the process of iterative mode updates. We find that smoothing the update with decay $\alpha_t < 1$ improves the stability of the algorithm:

$$\boldsymbol{\mu}_{t+1} = (1 - \alpha_t) \boldsymbol{\mu}_t + \alpha_t \boldsymbol{\mu}' \quad (21)$$

$$\text{where } \boldsymbol{\mu}' = \frac{1}{\sigma^2} \boldsymbol{\Sigma}_t \mathbf{W}_t^T (\mathbf{x} - \mathbf{b}_t). \quad (22)$$

Finally, by assuming the linear model $g_\theta(\mathbf{z}) \approx \mathbf{W}_T \mathbf{z} + \mathbf{b}_T$ at $\boldsymbol{\mu}_T$, the approximate Gaussian posterior is defined:

$$q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_T, \boldsymbol{\Sigma}_T), \quad (23)$$

$$\text{where } \boldsymbol{\Sigma}_T = \left(\frac{1}{\sigma^2} \mathbf{W}_T^T \mathbf{W}_T + \mathbf{I} \right)^{-1}. \quad (24)$$

We train our model by optimizing the ELBO in Eq.(4) by plugging in $q(\mathbf{z}|\mathbf{x})$ of Eq.(23). For sampling from the mul-

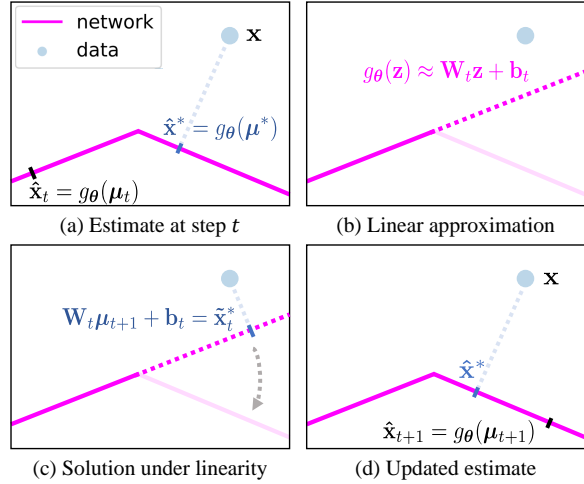


Figure 2. Illustration of the iterative update (Eq.(20))² for the posterior mode, drawn on the data space. (a) The posterior mode $\boldsymbol{\mu}^*$ corresponds to the best reconstruction $\hat{\mathbf{x}}^*$ of data \mathbf{x} on the network manifold, *i.e.* $\hat{\mathbf{x}}^* = g_\theta(\boldsymbol{\mu}^*)$. $\hat{\mathbf{x}}_t$ shows the estimate at step t . (b) We apply linear approximation to the network (dashed line). (c) We solve for $\boldsymbol{\mu}_{t+1}$ under this linear model (Eq.(20)). The network warps the result according to its manifold (dotted arrow). (d) Updated estimate. $\hat{\mathbf{x}}_{t+1}$ is now closer to $\hat{\mathbf{x}}^*$ than previous $\hat{\mathbf{x}}_t$.

tivariate Gaussian and propagating the gradient to the inference model, we calculate the Cholesky decomposition $\mathbf{L}\mathbf{L}^T = \boldsymbol{\Sigma}_T$ and apply the reparameterization $\mathbf{z} = \boldsymbol{\mu} + \mathbf{L}\boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is the standard Gaussian noise.

We highlight the notable characteristics of our approach:

- We iteratively update the posterior mode rather than solely relying on the amortized prediction. This is in spirit similar to semi-amortized inference (Kim et al., 2018; Marino et al., 2018; Krishnan et al., 2018), but critical differences are: (1) our method can make large jumps than prevalent gradient-based methods by exploiting the local linearity of the network, and (2) it is efficient and deterministic since it requires no sampling during updates. Fig. 3 intuitively depicts these effects.
- We gain the expressiveness of the full-covariance Gaussian posterior (Fig 4) where the covariance is analytically computed from the local behavior of the generative network. This is in contrast with normalizing flows (Rezende & Mohamed, 2015; Kingma et al., 2016; Tomczak & Welling, 2016) which introduce extra parameter overhead for the inference model and hence are prone to the amortization error.

²We here assume $\sigma^2 \rightarrow 0$ for the purpose of illustration. For $\sigma^2 > 0$, the prior shrinks the mode toward zero.

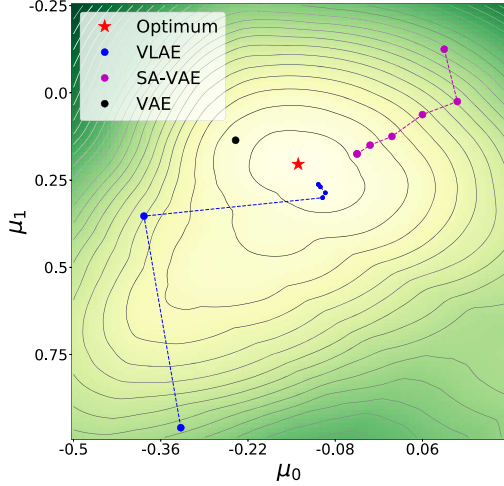


Figure 3. Illustration of the iterative mode update (Eq.(20)) on the ELBO landscape. The models are trained on MNIST using 2-dim latent variable. The update paths for the posterior mode μ are depicted for five steps. The VLAE obtains the closest estimate, making large jumps during the process for faster convergence.




Figure 4. The covariance matrix of $q(\mathbf{z}|\mathbf{x})$ from four different models. (a) VAE (Kingma & Welling, 2014). (b) Semi-Amortized VAE (Kim et al., 2018). (c) Householder Flow (Tomczak & Welling, 2016). (d) VLAE. While the former two only model diagonal elements, the latter two can model the full-covariance. The VLAE captures the richest correlations between dimensions.

3.4. Variational Laplace Autoencoders

So far, the proposed approach assumes the neural networks with rectified linear activations and Gaussian output. We here present a general framework *Variational Laplace Autoencoders* (VLAE), which are applicable to the general class of differentiable neural networks. To estimate the posterior $p_\theta(\mathbf{z}|\mathbf{x})$, VLAEs employ the *Laplace approximation* (Bishop, 2006) that finds a Gaussian approximation based on the local curvature at the posterior mode. We present how the Laplace method is incorporated for training deep generative models, and show the model in section 3.3 is a special case of VLAEs where it admits efficient computations.

Consider the problem of approximating the posterior $p_\theta(\mathbf{z}|\mathbf{x})$ where we only have access to the unnormalized density $p_\theta(\mathbf{x}, \mathbf{z})$. The Laplace method finds a Gaussian approximation $q(\mathbf{z}|\mathbf{x})$ centered on the mode of $p_\theta(\mathbf{z}|\mathbf{x})$ where the covariance is determined by the local curvature of

Algorithm 2 Variational Laplace Autoencoders

Input: generative model θ , inference model ϕ
 Sample $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$
 Initialize $\mu = \text{enc}_\phi(\mathbf{x})$
for $t = 0$ **to** $T - 1$ **do**
 Update μ (e.g. using gradient descent)
end for
 $\Sigma \leftarrow (-\nabla_{\mathbf{z}}^2 \ln p_\theta(\mathbf{x}, \mathbf{z})|_{\mathbf{z}=\mu})^{-1}$
 $q(\mathbf{z}|\mathbf{x}) \leftarrow \mathcal{N}(\mu, \Sigma)$
 Sample $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 Compute the Cholesky decomposition $\mathbf{L}\mathbf{L}^T = \Sigma$
 $\mathbf{z} \leftarrow \mu + \mathbf{L}\epsilon$
 Estimate the ELBO: $\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \ln p_\theta(\mathbf{x}, \mathbf{z}) - \ln q_\phi(\mathbf{z}|\mathbf{x})$
 Update generative model: $\theta \leftarrow \theta + \alpha \nabla_\theta \mathcal{L}_{\theta, \phi}(\mathbf{x})$
 Update inference model: $\phi \leftarrow \phi + \alpha \nabla_\phi \mathcal{L}_{\theta, \phi}(\mathbf{x})$

$\log p_\theta(\mathbf{x}, \mathbf{z})$. The overall procedure is largely divided into two parts: (1) finding the mode of posterior distribution and (2) computing the Gaussian approximation centered at the mode.

First, we iteratively search for the mode μ of $p_\theta(\mathbf{z}|\mathbf{x})$ via

$$\nabla_{\mathbf{z}} \log p_\theta(\mathbf{x}, \mathbf{z})|_{\mathbf{z}=\mu} = \mathbf{0}. \quad (25)$$

We can generally apply gradient-based optimization for this purpose. After determining the mode, we run the second-order Taylor expansion centered at μ :

$$\log p_\theta(\mathbf{x}, \mathbf{z}) \approx \log p_\theta(\mathbf{x}, \mu) - \frac{1}{2}(\mathbf{z} - \mu)^T \Lambda (\mathbf{z} - \mu), \quad (26)$$

$$\text{where } \Lambda = -\nabla_{\mathbf{z}}^2 \log p_\theta(\mathbf{x}, \mathbf{z})|_{\mathbf{z}=\mu}. \quad (27)$$

As this form is equivalent to Gaussian distribution, we define the approximate posterior as

$$q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mu, \Sigma), \quad (28)$$

$$\text{where } \Sigma^{-1} = \Lambda = -\nabla_{\mathbf{z}}^2 \log p_\theta(\mathbf{x}, \mathbf{z})|_{\mathbf{z}=\mu}. \quad (29)$$

This posterior distribution is then used to estimate the ELBO in Eq.(4) for training the generative model. Alg. 2 summarizes the proposed framework.

Gaussian output ReLU networks. We now make a connection to the approach in section 3.3. For the generative model defined in Eq.(18)–(19), the joint log-likelihood is

$$\begin{aligned} \log p_\theta(\mathbf{x}, \mathbf{z}) \\ = -\frac{1}{2\sigma^2}(\mathbf{x} - g_\theta(\mathbf{z}))^T(\mathbf{x} - g_\theta(\mathbf{z})) - \frac{1}{2}\mathbf{z}^T \mathbf{z} + C, \end{aligned} \quad (30)$$

where C is a constant independent of \mathbf{z} . Taking the gradient with respect to \mathbf{z} and setting it to zero,

$$\nabla_{\mathbf{z}} \log p_\theta(\mathbf{x}, \mathbf{z}) = -\frac{1}{\sigma^2} \frac{\partial g_\theta(\mathbf{z})^T}{\partial \mathbf{z}} (g_\theta(\mathbf{z}) - \mathbf{x}) - \mathbf{z} = \mathbf{0}.$$

By assuming linear approximation $g_\theta(\mathbf{z}) \approx \mathbf{W}_z \mathbf{z} + \mathbf{b}_z$ (Eq.(11)) and plugging it in, the solution is

$$\mathbf{z} = \frac{1}{\sigma^2} \left(\frac{1}{\sigma^2} \mathbf{W}_z^T \mathbf{W}_z + \mathbf{I} \right)^{-1} \mathbf{W}_z^T (\mathbf{x} - \mathbf{b}_z), \quad (31)$$

which is equivalent to the update equation of Eq.(20). Moreover, under the linear model above:

$$\mathbf{\Lambda} = -\nabla_{\mathbf{z}}^2 \log p(\mathbf{x}, \mathbf{z}) = \left(\frac{1}{\sigma^2} \mathbf{W}_z^T \mathbf{W}_z + \mathbf{I} \right), \quad (32)$$

which agree with Σ^{-1} in Eq.(7). That is, by using the local linearity assumption, we can estimate the covariance without the expensive computation of the Hessian of the generative network.

Bernoulli output ReLU networks. We illustrate an example for how VLAEs can be applied to Bernoulli output ReLU networks:

$$\log p_\theta(\mathbf{x}|\mathbf{z}) = \sum_i^n x_i \log y_i(\mathbf{z}) + (1 - x_i) \log(1 - y_i(\mathbf{z})),$$

where $\mathbf{y}(\mathbf{z}) = \frac{1}{1 + \exp(-g_\theta(\mathbf{z}))}$. (33)

Using $\nabla_{g_\theta(\mathbf{z})} \log p(\mathbf{x}|\mathbf{z}) = \mathbf{x} - \mathbf{y}(\mathbf{z})$ and chain rule, the gradient and the Hessian is

$$\nabla_{\mathbf{z}} \log p_\theta(\mathbf{x}, \mathbf{z}) = \frac{\partial g_\theta(\mathbf{z})^T}{\partial \mathbf{z}} (\mathbf{x} - \mathbf{y}(\mathbf{z})) - \mathbf{z}, \quad (34)$$

$$\begin{aligned} \nabla_{\mathbf{z}}^2 \log p_\theta(\mathbf{x}, \mathbf{z}) &= \frac{\partial^2 g_\theta(\mathbf{z})^T}{\partial \mathbf{z}^2} (\mathbf{x} - \mathbf{y}(\mathbf{z})) \\ &\quad - \frac{\partial g_\theta(\mathbf{z})^T}{\partial \mathbf{z}} \text{diag}(\mathbf{y}(\mathbf{z}) \cdot (\mathbf{1} - \mathbf{y}(\mathbf{z}))) \frac{\partial g_\theta(\mathbf{z})}{\partial \mathbf{z}} - \mathbf{I}. \end{aligned} \quad (35)$$

Plugging in the linear approximation $g_\theta(\mathbf{z}) \approx \mathbf{W}_z \mathbf{z} + \mathbf{b}_z$, the result simplifies to

$$\nabla_{\mathbf{z}} \log p_\theta(\mathbf{x}, \mathbf{z}) = \mathbf{W}_z^T (\mathbf{x} - \mathbf{y}(\mathbf{z})) - \mathbf{z}, \quad (36)$$

$$\nabla_{\mathbf{z}}^2 \log p_\theta(\mathbf{x}, \mathbf{z}) = -(\mathbf{W}_z^T \mathbf{S}_z \mathbf{W}_z + \mathbf{I}), \quad (37)$$

$$\text{where } \mathbf{S}_z = \text{diag}(\mathbf{y}(\mathbf{z}) \cdot (\mathbf{1} - \mathbf{y}(\mathbf{z}))). \quad (38)$$

To solve Eq.(36), we first apply the first-order approximation of $\mathbf{y}(\mathbf{z})$:

$$\mathbf{y}(\mathbf{z}') \approx \mathbf{y}(\mathbf{z}) + \frac{\partial \mathbf{y}(\mathbf{z})}{\partial \mathbf{z}} (\mathbf{z}' - \mathbf{z}) \quad (39)$$

$$= \mathbf{y}(\mathbf{z}) + \mathbf{S}_z \mathbf{W}_z (\mathbf{z}' - \mathbf{z}). \quad (40)$$

We plug it into Eq.(36) and solve the equation for zero,

$$\mathbf{z}' = (\mathbf{W}_z^T \mathbf{S}_z \mathbf{W}_z + \mathbf{I})^{-1} \mathbf{W}_z^T (\mathbf{x} - \mathbf{y}(\mathbf{z}) + \mathbf{S}_z \mathbf{W}_z \mathbf{z}).$$

This leads to the update equation for the mode similar to the Gaussian case (Eq.(20)):

$$\boldsymbol{\mu}_{t+1} = \Sigma_t \mathbf{W}_t^T (\mathbf{x} - \mathbf{b}_t), \quad (41)$$

$$\text{where } \Sigma_t = (\mathbf{W}_t^T \mathbf{S}_t \mathbf{W}_t + \mathbf{I})^{-1}, \mathbf{b}_t = (\mathbf{y}_t - \mathbf{S}_t \mathbf{W}_t \boldsymbol{\mu}_t).$$

After T updates, the approximate posterior distribution for the Bernoulli output distribution is defined as

$$q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_T, \Sigma_T). \quad (42)$$

3.5. Efficient Computation

For a network with width $O(D)$, the calculation of \mathbf{W}_z (Eq.(16)–(17)) and the update equation (Eq.(20)) requires a series of matrix-matrix multiplication and matrix inversion of $O(D^3)$ complexity, whereas the standard forward or backward propagation through the network takes a series of matrix-vector multiplication of $O(D^2)$ cost. As this computation can be burdensome for bigger networks³, we discuss efficient alternatives for: (1) iterative mode seeking and (2) covariance estimation.

Iterative mode seeking. For piece-wise linear networks, nonlinear variants of Conjugate Gradient (CG) can be effective as the CG solves a system of linear equations efficiently. It requires the evaluation of matrix-vector products $\mathbf{W}_z \mathbf{z}$ and $\mathbf{W}_z^T \mathbf{r}$, where \mathbf{r} is the residual $\mathbf{x} - (\mathbf{W}_z \mathbf{z} + \mathbf{b})$. These are computable during the forward and backward pass with $O(D^2)$ complexity per iteration. See Appendix for details.

For general differentiable neural networks, the gradient-based optimizers such as SGD, momentum or ADAM (Kingma & Ba, 2015) can be adopted to find the solution of Eq.(25). The backpropagation through the gradient-based updates requires the evaluation of Hessian-vector products but there are efficient approximations such as the finite differences (LeCun et al., 1993) used in Kim et al. (2018).

Covariance estimation. Instead of analytically calculating the precision matrix $\mathbf{\Lambda} = \Sigma^{-1} = \sigma^{-2} \mathbf{W}^T \mathbf{W} + \mathbf{I}$ (Eq.(32)) we may directly approximate the precision matrix using truncated SVD for top k singular values and vectors of $\mathbf{\Lambda}$. The truncated SVD can be iteratively performed using the power method where each iteration involves evaluation of the Jacobian-vector product $\mathbf{W}^T \mathbf{W}_z$. This can be computed through the forward and backward propagation at $O(D^2)$ cost. One way to further accelerate the convergence of the power iterations is to extend the amortized inference model to predict k vectors as seed vectors for the power method. However, we still need to compute the Cholesky decomposition of the covariance matrix for sampling.

4. Related Work

Cremer et al. (2018) and Krishnan et al. (2018) reveal that VAEs suffer from the inference gap between the ELBO and the marginal log-likelihood. Cremer et al. (2018) decompose this gap as the sum of the *approximation error* and

³In our experiments, the computational overhead is affordable with GPU acceleration. For example, our MNIST model takes 10 GPU hours for 1,000 epochs on one Titan X Pascal.

the *amortization error*. The approximation error results from the choice of a particular variational family, such as fully-factorized Gaussians, restricting the distribution to be factorial or more technically, have a diagonal covariance matrix. On the other hand, the amortization error is caused by the suboptimality of variational parameters due to the amortized predictions.

This work is most akin to the line of works that contribute to reducing the amortization error by iteratively updating the variational parameters to improve the approximate posterior (Kim et al., 2018; Marino et al., 2018; Krishnan et al., 2018). Distinct from the prior approaches which rely on gradient-based optimization, our method explores the local linearity of the network to make more efficient updates.

Regarding the approximation error, various approaches have been proposed for improved expressiveness of posterior approximation. Importance weighted autoencoders (Yuri et al., 2015) learn flexible posteriors using importance weighting. Tran et al. (2016) incorporate Gaussian processes to enrich posterior representation. Maaløe et al. (2016) augment the model with auxiliary variables and Salimans et al. (2015) use Markov chains with Hamiltonian dynamics. Normalizing flows (Tabak & Turner, 2013; Rezende & Mohamed, 2015; Kingma et al., 2016; Tomczak & Welling, 2016) transform a simple initial distribution to an increasingly flexible one by using a series of *flows*, invertible transformations whose determinant of the Jacobian is easy to compute. For example, the Householder flows (Tomczak & Welling, 2016) can represent a Gaussian distribution with a full covariance alike to our VLAEs. However, as opposed to the flow-based approaches, VLAEs introduce no additional parameters and are robust to the amortization error.

The Laplace approximation have been applied to estimate the uncertainty of weight parameters of neural networks (MacKay, 1992; Ritter et al., 2018). However, due to the high dimensionality of neural network parameters (often over millions), strong assumptions on the structure of the Hessian matrix are required to make the computation feasible (LeCun et al., 1990; Ritter et al., 2018). On the other hand, the latent dimension of deep generative models is typically in the hundreds, rendering our approach practical.

5. Experiments

We evaluate our approach on five popular datasets: MNIST (LeCun et al., 1998), Omniglot (Lake et al., 2013), Fashion-MNIST (Xiao et al., 2017), Street View House Numbers (SVHN) (Wang et al., 2011) and CIFAR-10 (Krizhevsky & Hinton, 2009). We verify the effectiveness of our approach not only on the ReLU networks with Gaussian output (section 5.1), but also on the Bernoulli output ReLU networks in section 3.4 on dynamically binarized MNIST (section 5.2).

We compare the VLAE with other recent VAE models, including (1) VAE (Kingma & Welling, 2014) using the standard fully-factorized Gaussian assumption, (2) Semi-Amortized VAE (SA-VAE) which extends the VAE using gradient-based updates of variational parameters (Kim et al., 2018; Marino et al., 2018; Krishnan et al., 2018) (3) VAE augmented with Householder Flow (VAE+HF) (Tomczak & Welling, 2016) (Rezende & Mohamed, 2015) which employs a series of Householder transformation to model the covariance of the Gaussian posterior. For bigger networks, we also include (4) VAE augmented with Inverse Autoregressive Flow (VAE+IAF) (Kingma et al., 2016).

We experiment two network settings: (1) a small network with one hidden layer. The latent variable dimension is 16 and the hidden layer dimension is 256. We double both dimensions for color datasets of SVHN and CIFAR10. (2) A bigger network with two hidden layers. The latent variable dimension is 50 and the hidden layer dimension is 500 for all datasets. For both settings, we apply ReLU activation to hidden layers and use the same architecture for the encoder and decoder. See Appendix for more experimental details.

The code is public at <http://vision.snu.ac.kr/projects/VLAE>.

5.1. Results of Gaussian Outputs

Table 1 summarizes the results on the Gaussian output ReLU networks in the small network setting. The VLAE outperforms other baselines with notable margins in all the datasets, proving the effectiveness of our approach. Remarkably, a single step of update ($T = 1$) leads to substantial improvement compared to the other models, and with more updates the performance further enhances. Fig. 5 depicts how data reconstructions improve with the update steps.

Table 2 shows the results with the bigger networks, which bring considerable improvements. The VLAE again attains the best results for all the datasets. The VAE+IAF is generally strong among the baselines whereas the SA-VAE is worse compared to others. One distinguishing trend compared to the small network results is that increasing the number of updates T often degrade the performance of the models. We suspect the increased depth due to the large number of updates or flows causes optimization difficulties, as we observe worse results on the training set as well.

For both network settings, the SA-VAE and VAE+HF show mixed results. We hypothesize the causes are as follows: (1) The SA-VAE update is noisy as the gradient of ELBO (Eq.(4)) is estimated using a single sample of \mathbf{z} . Hence, it may cause instability during training and may require more iterations than used in our experiments for better performance. (2) Although the VAE+HF is endowed with flexibility to represent correlations between the latent variable dimensions, it is prone to suffer from the amortization

Table 1. Log-likelihood results on small networks, estimated with 100 importance samples. Gaussian output is used except the last column with the Bernoulli output. T refers to the number of updates for the VLAE and SA-VAE (Kim et al., 2018) or the number of flows for the VAE+HF (Tomczak & Welling, 2016).

	MNIST	OMNI-GLOT	FASHION MNIST	SVHN	CIFAR10	BINARY MNIST
VAE	612.9	343.5	606.3	4555	2364	-96.73
SA-VAE						
$T=1$	614.1	341.4	606.7	4553	2366	-96.85
$T=2$	615.2	346.6	604.1	4551	2366	-96.73
$T=4$	612.8	348.6	606.6	4553	2366	-96.71
$T=8$	612.1	345.5	608.0	4559	2365	-96.89
VAE+HF						
$T=1$	610.5	341.5	604.3	4557	2366	-96.75
$T=2$	613.1	343.1	606.5	4569	2361	-96.52
$T=4$	612.9	333.8	604.9	4564	2362	-96.44
$T=8$	615.6	332.6	605.5	4536	2357	-96.14
VLAE						
$T=1$	638.6	362.0	614.9	4639	2374	-94.68
$T=2$	645.4	372.7	615.5	4681	2381	-94.46
$T=4$	649.9	372.3	615.6	4711	2387	-94.41
$T=8$	650.3	380.7	618.8	4718	2392	-94.57




Figure 5. Examples of images reconstructed by VLAE with five update iterations proceeding from left to right. The pixel correlations greatly improve on the first update. The ground-truth samples are shown on the rightmost.

error as it completely relies on the dynamic prediction of the inference network, whereas the inference problem becomes more complex with the enhanced flexibility.

On the other hand, we argue that the VLAE is able to make significant improvements in fewer steps because the VLAE update is more powerful and deterministic (Eq.(20)), thanks to the local linearity of ReLU networks. Moreover, the VLAE computes the covariance of latent variables directly from the generative model, thus providing greater expressiveness with less amortization error.

5.2. Results of Bernoulli Outputs

To show the generality of our approach on non-Gaussian output, we experiment ReLU networks with Bernoulli output on dynamically binarized MNIST. Each pixel is stochastically set to 1 or 0 with a probability proportional to the pixel

Table 2. Log-likelihood results on bigger networks, estimated with 5000 importance samples. See the caption of Table 1 for details.

	MNIST	OMNI-GLOT	FASHION MNIST	SVHN	CIFAR10	BINARY MNIST
VAE	1015	602.7	707.7	5162	2640	-85.38
SA-VAE						
$T=1$	984.7	598.5	706.4	5181	2639	-85.20
$T=2$	1006	589.8	708.3	5165	2639	-85.10
$T=4$	999.8	604.6	706.4	5172	2640	-85.43
$T=8$	990.0	602.7	697.0	5172	2639	-85.24
VAE+HF						
$T=1$	1028	602.5	715.6	5201	2637	-85.27
$T=2$	1020	603.1	710.4	5179	2636	-85.31
$T=4$	989.8	607.0	710.0	5209	2641	-85.22
$T=8$	944.3	608.5	714.6	5196	2640	-85.41
VAE+IAF						
$T=1$	1015	609.4	721.2	5037	2642	-84.26
$T=2$	1057	617.5	724.1	5150	2624	-84.16
$T=4$	1051	617.5	721.0	4994	2638	-84.03
$T=8$	1018	606.7	725.7	4951	2639	-83.80
VLAE						
$T=1$	1150	727.4	817.7	5324	2687	-83.72
$T=2$	1096	731.1	825.4	5159	2686	-83.84
$T=4$	1054	701.2	826.0	5231	2683	-83.73
$T=8$	1009	661.2	821.0	5341	2639	-83.60

intensity (Salakhutdinov & Murray, 2008).

The rightmost columns in Table 1–2 show the results on dynamically binarized MNIST. The VLAE attains the highest log-likelihood as in the Gaussian output experiments. The results demonstrate that VLAEs are also effective for non-Gaussian output models.

6. Conclusion

We presented *Variational Laplace Autoencoders* (VLAEs), which apply the Laplace approximation of the posterior for training deep generative models. The iterative mode updates and full-covariance Gaussian approximation using the curvature of the generative network enhances the expressive power of the posterior with less amortization error. The experiments demonstrated that on ReLU networks, the VLAEs outperformed other amortized or iterative models.

As future work, an important study may be to extend VLAEs to deep latent models based on the combination of top-down information and bottom-up inference (Kingma et al., 2016; Sønderby et al., 2016).

Acknowledgements

This work is supported by Samsung Advanced Institute of Technology, Korea-U.K. FP Programme through NRF of Korea (NRF-2017K1A3A1A16067245) and IITP grant funded by the Korea government (MSIP) (2019-0-01082).

References

- Bishop, C. M. *Pattern recognition and machine learning*. Springer-Verlag New York, Inc., Secaucus, NJ, 2006.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. In *CoNLL*, 2016.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A. C., and Bengio, Y. A recurrent latent variable model for sequential data. In *NeurIPS*, 2015.
- Cremer, C., Li, X., and Duvenaud, D. Inference suboptimality in variational autoencoders. In *ICML*, 2018.
- Goodfellow, I. J., Warde-farley, D., Mirza, M., Courville, A., and Bengio, Y. Maxout networks. In *ICML*, 2013.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D. J., and Wierstra, D. Draw: A recurrent neural network for image generation. In *ICML*, 2015.
- Gulrajani, I., Kumar, K., Faruk, A., Taiga, A. A., Visin, F., Vazquez, D., and Courville, A. C. Pixelvae: A latent variable model for natural images. In *ICLR*, 2017.
- Hinton, G. E. and Camp, D. V. Keeping the neural networks simple by minimizing the description length of the weights. In *COLT*, 1993.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Kim, Y., Wiseman, S., Millter, A. C., Sontag, D., and Rush, A. M. Semi-amortized variational autoencoders. In *ICML*, 2018.
- Kingma, D. and Ba, J. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *ICLR*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *NeurIPS*, 2016.
- Krishnan, R. G., Liang, D., and Hoffman, M. D. On the challenges of learning with inference networks on sparse high-dimensional data. In *AISTAT*, 2018.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, Computer Science Department, University of Toronto, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- Lake, B. M., Salakhutdinov, R. R., and Tenenbaum, J. One-shot learning by inverting a compositional causal process. In *NeurIPS*, 2013.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In *NeurIPS*, 1990.
- LeCun, Y., Simard, P. Y., and Pearlmutter, B. Automatic learning rate maximization by on-line estimation of the hessian's eigenvectors. In *NeurIPS*, 1993.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient based learning applied to document recognition. In *IEEE*, 1998.
- Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. Auxiliary deep generative models. In *ICML*, 2016.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, 2013.
- MacKay, D. J. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
- Marino, J., Yisong, Y., and Mandt, S. Iterative amortized inference. In *ICML*, 2018.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *ICML*, 2017.
- Montufar, G., Pascanu, R., Cho, K., and Bengio, Y. On the number of linear regions of deep neural networks. In *NeurIPS*, 2014.
- Pascanu, R., Montufar, G., and Bengio, Y. On the number of response regions of deep feedforward networks with piecewise linear activations. In *ICLR*, 2014.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. In *ICML*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- Ritter, H., Botev, A., and Barber, D. A scalable laplace approximation for neural networks. In *ICLR*, 2018.
- Roberts, A., Engel, J., Raffel, C., Hawthorne, C., and Eck, D. A hierarchical latent vector model for learning long-term structure in music. In *ICML*, 2018.
- Salakhutdinov, R. and Murray, I. On the quantitative analysis of deep belief networks. In *ICML*, 2008.

- Salimans, T., Kingma, D., and Welling, M. Markov chain monte carlo and variational inference: Bridging the gap. In *ICML*, 2015.
- Sønderby, C. K., Raiko, T. R., Maaløe, L., Sønderby, S. K., and Winther, O. How to train deep variational autoencoders and probabilistic ladder networks. In *ICML*, 2016.
- Tabak, E. G. and Turner, C. V. A family of nonparametric density estimation algorithms. In *Communications on Pure and Applied Mathematics*, 2013.
- Tipping, M. E. and Bishop, C. M. Probabilistic Principal Component Analysis. *J. R. Statist. Soc. B*, 61(3):611–622, 1999.
- Tomczak, J. M. and Welling, M. Improving variational autoencoders using householder flow. In *NeurIPS Workshop on Bayesian Deep Learning*, 2016.
- Tran, D., Ranganath, R., and Blei, D. M. The variational gaussian process. In *ICLR*, 2016.
- Wang, Y. N. T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NeurIPS*, 2011.
- Waterhouse, S. R., MacKay, D., and Robinson, A. J. Bayesian methods for mixtures of experts. In *NeurIPS*, 1996.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. In *arXiv*, 2017.
- Yuri, B., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. In *ICLR*, 2015.