
Semi-Amortized Variational Autoencoders

Yoon Kim¹ Sam Wiseman¹ Andrew C. Miller¹ David Sontag² Alexander M. Rush¹

Abstract

Amortized variational inference (AVI) replaces instance-specific local inference with a global inference network. While AVI has enabled efficient training of deep generative models such as variational autoencoders (VAE), recent empirical work suggests that inference networks can produce sub-optimal variational parameters. We propose a hybrid approach, to use AVI to initialize the variational parameters and run stochastic variational inference (SVI) to refine them. Crucially, the local SVI procedure is itself differentiable, so the inference network and generative model can be trained end-to-end with gradient-based optimization. This semi-amortized approach enables the use of rich generative models without experiencing the posterior-collapse phenomenon common in training VAEs for problems like text generation. Experiments show this approach outperforms strong autoregressive and variational baselines on standard text and image datasets.

1. Introduction

Variational inference (VI) (Jordan et al., 1999; Wainwright & Jordan, 2008) is a framework for approximating an intractable distribution by optimizing over a family of tractable surrogates. Traditional VI algorithms iterate over the observed data and update the variational parameters with closed-form coordinate ascent updates that exploit conditional conjugacy (Ghahramani & Beal, 2001). This style of optimization is challenging to extend to large datasets and non-conjugate models. However, recent advances in *stochastic* (Hoffman et al., 2013), *black-box* (Ranganath et al., 2014; 2016), and *amortized* (Mnih & Gregor, 2014; Kingma & Welling, 2014; Rezende et al., 2014) variational inference have made it possible to scale to large datasets and rich, non-conjugate models (see Blei et al. (2017), Zhang et al. (2017) for a review of modern methods).

¹School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA ²CSAIL & IMES, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence to: Yoon Kim <yoonkim@seas.harvard.edu>.

In *stochastic variational inference* (SVI), the variational parameters for each data point are randomly initialized and then optimized to maximize the evidence lower bound (ELBO) with, for example, gradient ascent. These updates are based on a subset of the data, making it possible to scale the approach. In *amortized variational inference* (AVI), the local variational parameters are instead predicted by an inference (or recognition) network, which is shared (i.e. amortized) across the dataset. Variational autoencoders (VAEs) are deep generative models that utilize AVI for inference and jointly train the generative model alongside the inference network.

SVI gives good local (i.e. instance-specific) distributions within the variational family but requires performing optimization for each data point. AVI has fast inference, but having the variational parameters be a parametric function of the input may be too strict of a restriction. As a secondary effect this may militate against learning a good generative model since its parameters may be updated based on suboptimal variational parameters. Cremer et al. (2018) observe that the *amortization gap* (the gap between the log-likelihood and the ELBO due to amortization) can be significant for VAEs, especially on complex datasets.

Recent work has targeted this amortization gap by combining amortized inference with iterative refinement during training (Hjelm et al., 2016; Krishnan et al., 2018). These methods use an encoder to initialize the local variational parameters, and then subsequently run an iterative procedure to refine them. To train with this hybrid approach, they utilize a separate training time objective. For example Hjelm et al. (2016) train the inference network to minimize the KL-divergence between the initial and the final variational distributions, while Krishnan et al. (2018) train the inference network with the usual ELBO objective based on the initial variational distribution.

In this work, we address the train/test objective mismatch and consider methods for training *semi-amortized variational autoencoders* (SA-VAE) in a fully end-to-end manner. We propose an approach that leverages differentiable optimization (Domke, 2012; Maclaurin et al., 2015; Belanger et al., 2017) and differentiates *through* SVI while training the inference network/generative model. We find that this method is able to both improve estimation of variational parameters and produce better generative models.

We apply our approach to train deep generative models of text and images, and observe that they outperform autoregressive/VAE/SVI baselines, in addition to direct baselines that combine VAE with SVI but do not perform end-to-end training. We also find that under our framework, we are able to utilize a powerful generative model without experiencing the “posterior-collapse” phenomenon often observed in VAEs, wherein the variational posterior collapses to the prior and the generative model ignores the latent variable (Bowman et al., 2016; Chen et al., 2017; Zhao et al., 2017). This problem has particularly made it very difficult to utilize VAEs for text, an important open issue in the field. With SA-VAE, we are able to outperform an LSTM language model by utilizing an LSTM generative model that maintains non-trivial latent representations. Code is available at <https://github.com/harvardnlp/sa-vae>.

2. Background

Notation Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a scalar valued function with partitioned inputs $\mathbf{u} = [\mathbf{u}_1, \dots, \mathbf{u}_m]$ such that $\sum_{i=1}^m \dim(\mathbf{u}_i) = n$. With a slight abuse of notation we define $f(\mathbf{u}_1, \dots, \mathbf{u}_m) = f([\mathbf{u}_1, \dots, \mathbf{u}_m])$. We denote $\nabla_{\mathbf{u}_i} f(\hat{\mathbf{u}}) \in \mathbb{R}^{\dim(\mathbf{u}_i)}$ to be the i -th block of the gradient of f evaluated at $\hat{\mathbf{u}} = [\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_m]$, and further use $\frac{df}{d\mathbf{v}}$ to denote the total derivative of f with respect to \mathbf{v} , which exists if \mathbf{u} is a differentiable function of \mathbf{v} . Note that in general $\nabla_{\mathbf{u}_i} f(\hat{\mathbf{u}}) \neq \frac{df}{d\mathbf{u}_i}$ since other components of \mathbf{u} could be a function of \mathbf{u}_i .¹ We also let $H_{\mathbf{u}_i, \mathbf{u}_j} f(\hat{\mathbf{u}}) \in \mathbb{R}^{\dim(\mathbf{u}_i) \times \dim(\mathbf{u}_j)}$ be the matrix formed by taking the i -th group of rows and the j -th group of columns of the Hessian of f evaluated at $\hat{\mathbf{u}}$. These definitions generalize straightforwardly when $f : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is a vector-valued function (e.g. $\frac{df}{d\mathbf{u}} \in \mathbb{R}^{n \times p}$).²

2.1. Variational Inference

Consider the following generative process for \mathbf{x} ,

$$\mathbf{z} \sim p(\mathbf{z}) \quad \mathbf{x} \sim p(\mathbf{x} | \mathbf{z}; \theta)$$

where $p(\mathbf{z})$ is the prior and $p(\mathbf{x} | \mathbf{z}; \theta)$ is given by a generative model with parameters θ . As maximizing the log-likelihood $\log p(\mathbf{x}; \theta) = \log \int_{\mathbf{z}} p(\mathbf{x} | \mathbf{z}; \theta) p(\mathbf{z}) d\mathbf{z}$ is usually intractable, variational inference instead defines a variational family of distributions $q(\mathbf{z}; \lambda)$ parameterized by λ and maximizes the evidence lower bound (ELBO)

$$\begin{aligned} \log p(\mathbf{x}; \theta) &\geq \mathbb{E}_{q(\mathbf{z}; \lambda)} [\log p(\mathbf{x} | \mathbf{z})] - \text{KL}[q(\mathbf{z}; \lambda) \| p(\mathbf{z})] \\ &= \text{ELBO}(\lambda, \theta, \mathbf{x}) \end{aligned}$$

The variational posterior, $q(\mathbf{z}; \lambda)$, is said to *collapse* to the prior if $\text{KL}[q(\mathbf{z}; \lambda) \| p(\mathbf{z})] \approx 0$. In the general case we are

¹This will indeed be the case in our approach: when we calculate $\text{ELBO}(\lambda_K, \theta, \mathbf{x})$, λ_K is a function of the data point \mathbf{x} , the generative model θ , and the inference network ϕ (Section 3).

²Total derivatives/Jacobians are usually denoted with row vectors but we denote them with column vectors for clearer notation.

given a dataset $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$ and need to find variational parameters $\lambda^{(1)}, \dots, \lambda^{(N)}$ and generative model parameters θ that jointly maximize $\sum_{i=1}^N \text{ELBO}(\lambda^{(i)}, \theta, \mathbf{x}^{(i)})$.

2.2. Stochastic Variational Inference

We can apply SVI (Hoffman et al., 2013) with gradient ascent to approximately maximize the above objective:³

1. Sample $\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})$
2. Randomly initialize λ_0
3. For $k = 0, \dots, K - 1$, set $\lambda_{k+1} = \lambda_k + \alpha \nabla_{\lambda} \text{ELBO}(\lambda_k, \theta, \mathbf{x})$
4. Update θ based on $\nabla_{\theta} \text{ELBO}(\lambda_K, \theta, \mathbf{x})$

Here K is the number of SVI iterations and α is the learning rate. (Note that θ is updated based on the gradient $\nabla_{\theta} \text{ELBO}(\lambda_K, \theta, \mathbf{x})$ and not the total derivative $\frac{d \text{ELBO}(\lambda_K, \theta, \mathbf{x})}{d\theta}$. The latter would take into account the fact that λ_k is a function of θ for $k > 0$.)

SVI optimizes directly for instance-specific variational distributions, but may require running iterative inference for a large number of steps. Further, because of this block coordinate ascent approach the variational parameters λ are optimized separately from θ , potentially making it difficult for θ to adapt to local optima.

2.3. Amortized Variational Inference

AVI uses a global parametric model to predict the local variational parameters for each data point. A particularly popular application of AVI is in training the variational autoencoder (VAE) (Kingma & Welling, 2014), which runs an inference network (i.e. encoder) $\text{enc}(\cdot)$ parameterized by ϕ over the input to obtain the variational parameters:

1. Sample $\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})$
2. Set $\lambda = \text{enc}(\mathbf{x}; \phi)$
3. Update θ based on $\nabla_{\theta} \text{ELBO}(\lambda, \theta, \mathbf{x})$ (which in this case is equal to the total derivative)
4. Update ϕ based on

$$\frac{d \text{ELBO}(\lambda, \theta, \mathbf{x})}{d\phi} = \frac{d\lambda}{d\phi} \nabla_{\lambda} \text{ELBO}(\lambda, \theta, \mathbf{x})$$

The inference network is learned jointly alongside the generative model with the same loss function, allowing the pair to coadapt. Additionally inference for AVI involves running the inference network over the input, which is usually much faster than running iterative optimization on the ELBO. Despite these benefits, requiring the variational parameters to be a parametric function of the input may be too strict of a

³While we describe the various algorithms for a specific data point, in practice we use mini-batches.

restriction and can lead to an amortization gap. This gap can propagate forward to hinder the learning of the generative model if θ is updated based on suboptimal λ .

3. Semi-Amortized Variational Autoencoders

Semi-amortized variational autoencoders (SA-VAE) utilize an inference network over the input to give the initial variational parameters, and subsequently run SVI to refine them. One might appeal to the universal approximation theorem (Hornik et al., 1989) and question the necessity of additional SVI steps given a rich-enough inference network. However, in practice we find that the variational parameters found from VAE are usually not optimal even with a powerful inference network, and the amortization gap can be significant especially on complex datasets (Cremer et al., 2018; Krishnan et al., 2018).

SA-VAE models are trained using a combination of AVI and SVI steps:

1. Sample $\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})$
2. Set $\lambda_0 = \text{enc}(\mathbf{x}; \phi)$
3. For $k = 0, \dots, K - 1$, set

$$\lambda_{k+1} = \lambda_k + \alpha \nabla_{\lambda} \text{ELBO}(\lambda_k, \theta, \mathbf{x})$$
4. Update θ based on $\frac{d \text{ELBO}(\lambda_K, \theta, \mathbf{x})}{d\theta}$
5. Update ϕ based on $\frac{d \text{ELBO}(\lambda_K, \theta, \mathbf{x})}{d\phi}$

Note that for training we need to compute the total derivative of the final ELBO with respect to θ, ϕ (i.e. steps 4 and 5 above). Unlike with AVI, in order to update the encoder and generative model parameters, this total derivative requires backpropagating through the SVI updates. Specifically this requires backpropagating through gradient ascent (Domke, 2012; Maclaurin et al., 2015).

Following past work, this backpropagation step can be done efficiently with fast Hessian-vector products (LeCun et al., 1993; Pearlmutter, 1994). In particular, consider the case where we perform one step of refinement, $\lambda_1 = \lambda_0 + \alpha \nabla_{\lambda} \text{ELBO}(\lambda_0, \theta, \mathbf{x})$, and for brevity let $\mathcal{L} = \text{ELBO}(\lambda_1, \theta, \mathbf{x})$. To backpropagate through this, we receive the derivative $\frac{d\mathcal{L}}{d\lambda_1}$ and use the chain rule,

$$\begin{aligned} \frac{d\mathcal{L}}{d\lambda_0} &= \frac{d\lambda_1}{d\lambda_0} \frac{d\mathcal{L}}{d\lambda_1} = (\mathbf{I} + \alpha \mathbf{H}_{\lambda, \lambda} \text{ELBO}(\lambda_0, \theta, \mathbf{x})) \frac{d\mathcal{L}}{d\lambda_1} \\ &= \frac{d\mathcal{L}}{d\lambda_1} + \alpha \mathbf{H}_{\lambda, \lambda} \text{ELBO}(\lambda_0, \theta, \mathbf{x}) \frac{d\mathcal{L}}{d\lambda_1} \end{aligned}$$

We can then backpropagate $\frac{d\mathcal{L}}{d\lambda_0}$ through the inference network to calculate the total derivative, i.e. $\frac{d\mathcal{L}}{d\phi} = \frac{d\lambda_0}{d\phi} \frac{d\mathcal{L}}{d\lambda_0}$. Similar rules can be used to derive $\frac{d\mathcal{L}}{d\theta}$.⁴ The full forward/backward step, which uses gradient descent with momentum on the negative ELBO, is shown in Algorithm 1.

⁴We refer the reader to Domke (2012) for the full derivation.

Algorithm 1 Semi-Amortized Variational Autoencoders

Input: inference network ϕ , generative model θ ,
 inference steps K , learning rate α , momentum γ ,
 loss function $f(\lambda, \theta, \mathbf{x}) = -\text{ELBO}(\lambda, \theta, \mathbf{x})$

Sample $\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})$
 $\lambda_0 \leftarrow \text{enc}(\mathbf{x}; \phi)$
 $v_0 \leftarrow 0$

for $k = 0$ **to** $K - 1$ **do**
 $v_{k+1} \leftarrow \gamma v_k - \nabla_{\lambda} f(\lambda_k, \theta, \mathbf{x})$
 $\lambda_{k+1} \leftarrow \lambda_k + \alpha v_{k+1}$
end for

$\mathcal{L} \leftarrow f(\lambda_K, \theta, \mathbf{x})$
 $\bar{\lambda}_K \leftarrow \nabla_{\lambda} f(\lambda_K, \theta, \mathbf{x})$
 $\bar{\theta} \leftarrow \nabla_{\theta} f(\lambda_K, \theta, \mathbf{x})$
 $\bar{v}_K \leftarrow 0$

for $k = K - 1$ **to** 0 **do**
 $\bar{v}_{k+1} \leftarrow \bar{v}_{k+1} + \alpha \bar{\lambda}_{k+1}$
 $\bar{\lambda}_k \leftarrow \bar{\lambda}_{k+1} - \mathbf{H}_{\lambda, \lambda} f(\lambda_k, \theta, \mathbf{x}) \bar{v}_{k+1}$
 $\bar{\theta} \leftarrow \bar{\theta} - \mathbf{H}_{\theta, \lambda} f(\lambda_k, \theta, \mathbf{x}) \bar{v}_{k+1}$
 $\bar{v}_k \leftarrow \gamma \bar{v}_{k+1}$
end for

$\frac{d\mathcal{L}}{d\theta} \leftarrow \bar{\theta}$
 $\frac{d\mathcal{L}}{d\phi} \leftarrow \frac{d\lambda_0}{d\phi} \bar{\lambda}_0$
 Update θ, ϕ based on $\frac{d\mathcal{L}}{d\theta}, \frac{d\mathcal{L}}{d\phi}$

In our implementation we calculate Hessian-vector products with finite differences (LeCun et al., 1993; Domke, 2012), which was found to be more memory-efficient than automatic differentiation (and therefore crucial for scaling our approach to rich inference networks/generative models). Specifically, we estimate $\mathbf{H}_{\mathbf{u}_i, \mathbf{u}_j} f(\hat{\mathbf{u}}) \mathbf{v}$ with

$$\begin{aligned} \mathbf{H}_{\mathbf{u}_i, \mathbf{u}_j} f(\hat{\mathbf{u}}) \mathbf{v} &\approx \frac{1}{\epsilon} \left(\nabla_{\mathbf{u}_i} f(\hat{\mathbf{u}}_0, \dots, \hat{\mathbf{u}}_j + \epsilon \mathbf{v}, \dots, \hat{\mathbf{u}}_m) \right. \\ &\quad \left. - \nabla_{\mathbf{u}_i} f(\hat{\mathbf{u}}_0, \dots, \hat{\mathbf{u}}_j, \dots, \hat{\mathbf{u}}_m) \right) \end{aligned}$$

where ϵ is some small number (we use $\epsilon = 10^{-5}$).⁵ We further clip the results (i.e. rescale the results if the norm exceeds a threshold) before and after each Hessian-vector product as well as during SVI, which helped mitigate exploding gradients and further gave better training signal to the inference network.⁶ See Appendix A for details.

⁵Since in our case the ELBO is a non-deterministic function due to sampling (and dropout, if applicable), care must be taken when calculating Hessian-vector product with finite differences to ensure that the source of randomness is the same when calculating the two gradient expressions.

⁶Without gradient clipping, in addition to numerical issues we empirically observed the model to degenerate to a case whereby it learned to rely too much on iterative inference, and thus the initial parameters from the inference network were poor. Another way to provide better signal to the inference network is to train against a weighted sum $\sum_{k=0}^K w_k \text{ELBO}(\lambda_k, \theta, \mathbf{x})$ for $w_k \geq 0$.

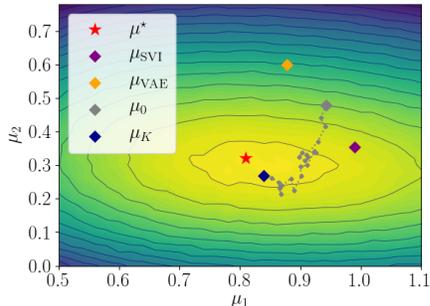


Figure 1. ELBO landscape with the oracle generative model as a function of the variational posterior means μ_1, μ_2 for a randomly chosen test point. Variational parameters obtained from VAE, SVI are shown as $\mu_{\text{VAE}}, \mu_{\text{SVI}}$ and the initial/final parameters from SA-VAE are shown as μ_0 and μ_K (along with the intermediate points). SVI/SA-VAE are run for 20 iterations. The optimal point, found from grid search, is shown as μ^* .

4. Experiments

We apply our approach to train generative models on a synthetic dataset in addition to text/images. For all experiments we utilize stochastic gradient descent with momentum on the negative ELBO. Our prior is the spherical Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and the variational posterior is diagonal Gaussian, where the variational parameters are given by the mean vector and the diagonal log variance vector, i.e. $\lambda = [\mu, \log \sigma^2]$.

In preliminary experiments we also experimented with natural gradients, other optimization algorithms, and learning the learning rates, but found that these did not significantly improve results. Full details regarding hyperparameters/model architectures for all experiments are in Appendix B.

4.1. Synthetic Data

We first apply our approach to a synthetic dataset where we have access to the true underlying generative model of discrete sequences. We generate synthetic sequential data according to the following oracle generative process with 2-dimensional latent variables and \mathbf{x}_t :

$$z_1, z_2 \sim \mathcal{N}(0, 1) \quad \mathbf{h}_t = \text{LSTM}(\mathbf{h}_{t-1}, \mathbf{x}_t) \\ \mathbf{x}_{t+1} \sim \text{softmax}(\text{MLP}([\mathbf{h}_t, z_1, z_2]))$$

We initialize the LSTM/MLP randomly as θ , where the LSTM has a single layer with hidden state/input dimension equal to 100. We generate for 5 time steps (so each example is given by $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_5]$) with a vocabulary size of 1000 for each \mathbf{x}_t . Training set consists of 5000 points. See Appendix B.1 for the exact setup.

We fix this oracle generative model $p(\mathbf{x} | \mathbf{z}; \theta)$ and learn an inference network (also a one-layer LSTM) with VAE and SA-VAE.⁷ For a randomly selected test point, we plot the

⁷With a fixed oracle, these models are technically not VAEs as VAE usually implies that the the generative model is learned (alongside the encoder).

MODEL	ORACLE GEN	LEARNED GEN
VAE	≤ 21.77	≤ 27.06
SVI	≤ 22.33	≤ 25.82
SA-VAE	≤ 20.13	≤ 25.21
TRUE NLL (EST)	19.63	—

Table 1. Variational upper bounds for the various models on the synthetic dataset, where SVI/SA-VAE is trained/tested with 20 steps. TRUE NLL (EST) is an estimate of the true negative log-likelihood (i.e. entropy of the data-generating distribution) estimated with 1000 samples from the prior. ORACLE GEN uses the oracle generative model and LEARNED GEN learns the generative network.

ELBO landscape in Figure 1 as a function of the variational posterior means (μ_1, μ_2) learned from the different methods. For SVI/SA-VAE we run iterative optimization for 20 steps. Finally we also show the optimal variational parameters found from grid search.

As can be seen from Figure 1, the variational parameters from running SA-VAE are closest to the optimum while those obtained from SVI and VAE are slightly further away. In Table 1 we show the variational upper bounds (i.e. negative ELBO) on the negative log-likelihood (NLL) from training the various models with both the oracle/learned generative model, and find that SA-VAE outperforms VAE/SVI in both cases.

4.2. Text

The next set of experiments is focused on text modeling on the Yahoo questions corpus from Yang et al. (2017). Text modeling with deep generative models has been a challenging problem, and few approaches have been shown to produce rich generative models that do not collapse to standard language models. Ideally a deep generative model trained with variational inference would make use of the latent space (i.e. maintain a nonzero KL term) while accurately modeling the underlying distribution.

Our architecture and hyperparameters are identical to the LSTM-VAE baselines considered in Yang et al. (2017), except that we train with SGD instead of Adam, which was found to perform better for training LSTMs. Specifically, both the inference network and the generative model are one-layer LSTMs with 1024 hidden units and 512-dimensional word embeddings. The last hidden state of the encoder is used to predict the vector of variational posterior means/log variances. The sample from the variational posterior is used to predict the initial hidden state of the generative LSTM and additionally fed as input at each time step. The latent variable is 32-dimensional. Following previous works (Bowman et al., 2016; Sønderby et al., 2016; Yang et al., 2017), for all the variational models we utilize a KL-cost annealing strategy whereby the multiplier on the KL term is increased linearly from 0.1 to 1.0 each batch over 10 epochs. Appendix B.2 has the full architecture/hyperparameters.

MODEL	NLL	KL	PPL
LSTM-LM	334.9	—	66.2
LSTM-VAE	⩽ 342.1	0.0	⩽ 72.5
LSTM-VAE + INIT	⩽ 339.2	0.0	⩽ 69.9
CNN-LM	335.4	—	66.6
CNN-VAE	⩽ 333.9	6.7	⩽ 65.4
CNN-VAE + INIT	⩽ 332.1	10.0	⩽ 63.9
LM	329.1	—	61.6
VAE	⩽ 330.2	0.01	⩽ 62.5
VAE + INIT	⩽ 330.5	0.37	⩽ 62.7
VAE + WORD-DROP 25%	⩽ 334.2	1.44	⩽ 65.6
VAE + WORD-DROP 50%	⩽ 345.0	5.29	⩽ 75.2
SVI ($K = 10$)	⩽ 331.4	0.16	⩽ 63.4
SVI ($K = 20$)	⩽ 330.8	0.41	⩽ 62.9
SVI ($K = 40$)	⩽ 329.8	1.01	⩽ 62.2
VAE + SVI ($K = 10$)	⩽ 331.2	7.85	⩽ 63.3
VAE + SVI ($K = 20$)	⩽ 330.5	7.80	⩽ 62.7
VAE + SVI + KL ($K = 10$)	⩽ 330.3	7.95	⩽ 62.5
VAE + SVI + KL ($K = 20$)	⩽ 330.1	7.81	⩽ 62.3
SA-VAE ($K = 10$)	⩽ 327.6	5.13	⩽ 60.5
SA-VAE ($K = 20$)	⩽ 327.5	7.19	⩽ 60.4

Table 2. Results on text modeling on the Yahoo dataset. Top results are from Yang et al. (2017), while the bottom results are from this work (+ INIT means the encoder is initialized with a pretrained language model, while models with + WORD-DROP are trained with word-dropout). NLL/KL numbers are averaged across examples, and PPL refers to perplexity. K refers to the number of inference steps used for training/testing.

In addition to autoregressive/VAE/SVI baselines, we consider two other approaches that also combine amortized inference with iterative refinement. The first approach is from Krishnan et al. (2018), where the generative model takes a gradient step based on the final variational parameters and the inference network takes a gradient step based on the initial variational parameters, i.e. we update θ based on $\nabla_{\theta} \text{ELBO}(\lambda_K, \theta, \mathbf{x})$ and update ϕ based on $\frac{d\lambda_0}{d\phi} \nabla_{\lambda} \text{ELBO}(\lambda_0, \theta, \mathbf{x})$. The forward step (steps 1-3 in Section 3) is identical to SA-VAE. We refer to this baseline as VAE + SVI.

In the second approach, based on Salakhutdinov & Larochelle (2010) and Hjelm et al. (2016), we train the inference network to minimize the KL-divergence between the initial and the final variational distributions, keeping the latter fixed. Specifically, letting $g(\nu, \omega) = \text{KL}[q(\mathbf{z}; \nu) \| q(\mathbf{z}; \omega)]$, we update θ based on $\nabla_{\theta} \text{ELBO}(\lambda_K, \theta, \mathbf{x})$ and update ϕ based on $\frac{d\lambda_0}{d\phi} \nabla_{\nu} g(\lambda_0, \lambda_K)$. Note that the inference network is not updated based on $\frac{dg}{d\phi}$, which would take into account the fact that both λ_0 and λ_K are functions of ϕ . We found $g(\lambda_0, \lambda_K)$ to perform better than the reverse direction $g(\lambda_K, \lambda_0)$. We refer to this setup as VAE + SVI + KL.

Results from the various models are shown in Table 2. Our baseline models (LM/VAE/SVI in Table 2) are already quite strong and outperform the models considered in Yang et al. (2017). However models trained with VAE/SVI make neg-

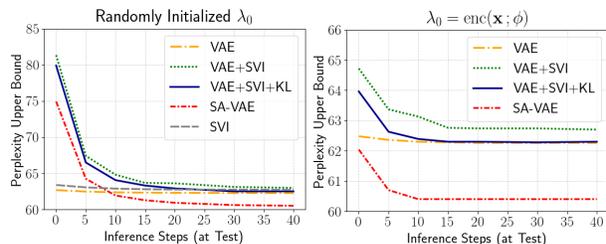


Figure 2. (Left) Perplexity upper bound of various models when trained with 20 steps (except for VAE) and tested with varying number of SVI steps from random initialization. (Right) Same as the left except that SVI is initialized with variational parameters obtained from the inference network.

ligible use of the latent variable and practically collapse to a language model, negating the benefits of using latent variables.⁸ In contrast, models that combine amortized inference with iterative refinement make use of the latent space and the KL term is significantly above zero.⁹ VAE + SVI and VAE + SVI + KL do not outperform a language model, and while SA-VAE only modestly outperforms it, to our knowledge this is one of the first instances in which we are able to train an LSTM generative model that does not ignore the latent code and outperforms a language model.

One might wonder if the improvements are coming from simply having a more flexible inference scheme at test time, rather than from learning a better generative model. To test this, for the various models we discard the inference network at test time and perform SVI for a variable number of steps from random initialization. The results are shown in Figure 2 (left). It is clear that the learned generative model (and the associated ELBO landscape) is quite different—it is not possible to train with VAE and perform SVI at test time to obtain the same performance as SA-VAE (although the performance of VAE does improve slightly from 62.7 to 62.3 when we run SVI for 40 steps from random initialization).

Figure 2 (right) has the results for a similar experiment where we refine the variational parameters initialized from the inference network for a variable number of steps at test time. We find that the inference network provides better initial parameters than random initialization and thus requires fewer iterations of SVI to reach the optimum. We do not observe improvements for running more refinement steps than was used in training at test time. Interestingly, SA-VAE without any refinement steps at test time has a substantially nonzero KL term (KL = 6.65, PPL = 62.0). This indicates that the posterior-collapse phenomenon when

⁸Models trained with word dropout (+ WORD-DROP in Table 2) do make use of the latent space but significantly underperform a language model.

⁹A high KL term does not necessarily imply that the latent variable is being utilized in a meaningful way (it could simply be due to bad optimization). In Section 5.1 we investigate the learned latent space in more detail.

MODEL	NLL
IWAE (Burda et al., 2015a)	103.38
LADDER VAE (Sønderby et al., 2016)	102.11
RBM (Burda et al., 2015b)	100.46
DISCRETE VAE (Rolfe, 2017)	97.43
DRAW (Gregor et al., 2015)	≤ 96.50
CONV DRAW (Gregor et al., 2016)	≤ 91.00
VLAE (Chen et al., 2017)	89.83
VAMPRIOR (Tomczak & Welling, 2018)	89.76
<hr/>	
GATED PIXELCNN	90.59
VAE	≤ 90.43 (0.98)
SVI ($K = 10$)	≤ 90.65 (0.02)
SVI ($K = 20$)	≤ 90.51 (0.06)
SVI ($K = 40$)	≤ 90.44 (0.27)
SVI ($K = 80$)	≤ 90.27 (1.65)
VAE + SVI ($K = 10$)	≤ 90.26 (1.69)
VAE + SVI ($K = 20$)	≤ 90.19 (2.40)
VAE + SVI + KL ($K = 10$)	≤ 90.24 (2.42)
VAE + SVI + KL ($K = 20$)	≤ 90.21 (2.83)
SA-VAE ($K = 10$)	≤ 90.20 (1.83)
SA-VAE ($K = 20$)	≤ 90.05 (2.78)

Table 3. Results on image modeling on the OMNIGLOT dataset. Top results are from prior works, while the bottom results are from this work. GATED PIXELCNN is our autoregressive baseline, and K refers to the number of inference steps during training/testing. For the variational models the KL portion of the ELBO is shown in parentheses.

training LSTM-based VAEs for text is partially due to optimization issues. Finally, while Yang et al. (2017) found that initializing the encoder with a pretrained language model improved performance (+ INIT in Table 2), we did not observe this on our baseline VAE model when we trained with SGD and hence did not pursue this further.

4.3. Images

We next apply our approach to model images on the OMNIGLOT dataset (Lake et al., 2015).¹⁰ While posterior collapse is less of an issue for VAEs trained on images, we still expect that improving the amortization gap would result in generative models that better model the underlying data and make more use of the latent space. We use a three-layer ResNet (He et al., 2016) as our inference network. The generative model first transforms the 32-dimensional latent vector to the image spatial resolution, which is concatenated with the original image and fed to a 12-layer Gated PixelCNN (van den Oord et al., 2016) with varying filter sizes, followed by a final sigmoid layer. We employ the same KL-cost annealing schedule as in the text experiments. See Appendix B.3 for the exact architecture/hyperparameters.

Results from the various models are shown in Table 3. Our findings are largely consistent with results from text: the semi-amortized approaches outperform VAE/SVI baselines, and further they learn generative models that make more

¹⁰We focus on the more complex OMNIGLOT dataset instead of the simpler MNIST dataset as prior work has shown that the amortization gap on MNIST is minimal (Cremer et al., 2018).

use of the latent representations (i.e. KL portion of the loss is higher). Even with 80 steps of SVI we are unable to perform as well as SA-VAE trained with 10 refinement steps, indicating the importance of good initial parameters provided by the inference network. In Appendix C we further investigate the performance of VAE and SA-VAE as we vary the training set size and the capacity of the inference network/generative model. We find that SA-VAE outperforms VAE and has higher latent variable usage in all scenarios. We note that we do not outperform the state-of-the-art models that employ hierarchical latent variables and/or more sophisticated priors (Chen et al., 2017; Tomczak & Welling, 2018). However these additions are largely orthogonal to our approach and we hypothesize they will also benefit from combining amortized inference with iterative refinement.¹¹

5. Discussion

5.1. Learned Latent Space

For the text model we investigate what the latent variables are learning through saliency analysis with our best model (SA-VAE trained with 20 steps). Specifically, we calculate the output saliency of each token \mathbf{x}_t with respect to \mathbf{z} as

$$\mathbb{E}_{q(\mathbf{z};\lambda)} \left[\left\| \frac{d \log p(\mathbf{x}_t | \mathbf{x}_{<t}, \mathbf{z}; \theta)}{d\mathbf{z}} \right\|_2 \right]$$

where $\|\cdot\|_2$ is the l_2 norm and the expectation is approximated with 5 samples from the variational posterior. Saliency is therefore a measure of how much the latent variable is being used to predict a particular token.

We visualize the saliency of a few examples from the test set in Figure 3 (top). Each example consists of a question followed by an answer from the Yahoo corpus. From a qualitative analysis several things are apparent: the latent variable seems to encode question type (i.e. if, what, how, why, etc.) and therefore saliency is high for the first word; content words (nouns, adjectives, lexical verbs) have much higher saliency than function words (determiners, prepositions, conjunctions, etc.); saliency of the $\langle /s \rangle$ token is quite high, indicating that the length information is also encoded in the latent space. In the third example we observe that the left parenthesis has higher saliency than the right parenthesis (0.32 vs. 0.24 on average across the test set), as the latter can be predicted by conditioning on the former rather than on the latent representation \mathbf{z} .

The previous definition of saliency measures the influence of \mathbf{z} on the *output* \mathbf{x}_t . We can also roughly measure the influence of the *input* \mathbf{x}_t on the latent representation \mathbf{z} , which we refer to as input saliency:

$$\left\| \mathbb{E}_{q(\mathbf{z};\lambda)} \left[\frac{d\|\mathbf{z}\|_2}{d\mathbf{w}_t} \right] \right\|_2$$

¹¹Indeed, Cremer et al. (2018) observe that the amortization gap can be substantial for VAE trained with richer variational families.

Semi-Amortized Variational Autoencoders



Figure 3. (Top) Saliency visualization of some examples from the test set. Here the saliency values are rescaled to be between 0-100 within each example for easier visualization. Red indicates higher saliency values. (Middle) Input saliency of the first test example from the top (in blue), in addition to two sample outputs generated from the variational posterior (with their saliency values in red). (Bottom) Same as the middle except we use a made-up example. Best viewed in color.

Here w_t is the encoder word embedding for x_t .¹² We visualize the input saliency for a test example (Figure 3, middle) and a made-up example (Figure 3, bottom). Under each input example we also visualize a two samples from the variational posterior, and find that the generated examples are often meaningfully related to the input example.¹³

We quantitatively analyze output saliency across part-of-speech, token position, word frequency, and log-likelihood in Figure 4: nouns (NN), adjectives (JJ), verbs (VB), numbers (CD), and the $\langle s \rangle$ token have higher saliency than conjunctions (CC), determiners (DT), prepositions (IN), and the TO token—the latter are relatively easier to predict by conditioning on previous tokens; similarly, on average, tokens occurring earlier have much higher saliency than those

¹²As the norm of \mathbf{z} is a rather crude measure, a better measure would be obtained by analyzing the spectra of the Jacobian $\frac{d\mathbf{z}}{dw_t}$. However this is computationally too expensive to calculate for each token in the corpus.

¹³We first sample $\mathbf{z} \sim q(\mathbf{z}; \lambda_K)$ then $\mathbf{x} \sim p(\mathbf{x} | \mathbf{z}; \theta)$. When sampling $\mathbf{x}_t \sim p(\mathbf{x}_t | \mathbf{x}_{<t}, \mathbf{z})$ we sample with temperature $T = 0.25$, i.e. $p(\mathbf{x}_t | \mathbf{x}_{<t}, \mathbf{z}) = \text{softmax}(\frac{1}{T}\mathbf{s}_t)$ where \mathbf{s}_t is the vector with scores for all words. We found the generated examples to be related to the original (in some way) in roughly half the cases.

occurring later (Figure 4 shows absolute position but the plot is similar with relative position); the latent variable is used much more when predicting rare tokens; there is some negative correlation between saliency and log-likelihood (-0.51), though this relationship does not always hold—e.g. $\langle s \rangle$ has high saliency but is relatively easy to predict with an average log-likelihood of -1.61 (vs. average log-likelihood of -4.10 for all tokens). Appendix D has the corresponding analysis for input saliency, which are qualitatively similar.

These results seem to suggest that the latent variables are encoding interesting and potentially interpretable aspects of language. While left as future work, it is possible that manipulations in the latent space of a model learned this way could lead to controlled generation/manipulation of output text (Hu et al., 2017; Mueller et al., 2017).

5.2. Limitations

A drawback of our approach (and other non-amortized inference methods) is that each training step requires backpropagating through the generative model multiple times, which can be costly especially if the generative model is expensive to compute (e.g. LSTM/PixelCNN). This may potentially be mitigated through more sophisticated *meta learning* ap-

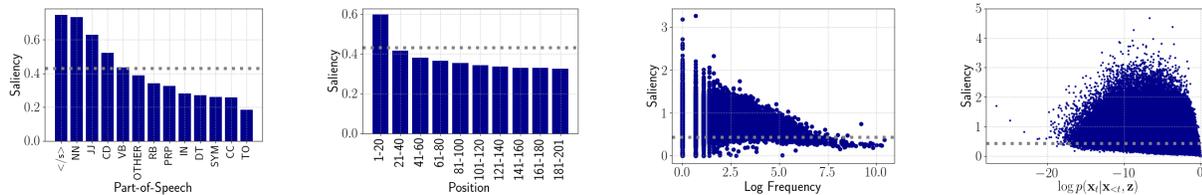


Figure 4. Output saliency by part-of-speech tag, position, log frequency, and log-likelihood. See Section 5.1 for the definitions of output saliency. The dotted gray line in each plot shows the average saliency across all words.

proaches (Andrychowicz et al., 2016; Marino et al., 2018), or with more efficient use of the past gradient information during SVI via averaging (Schmidt et al., 2013) or importance sampling (Sakaya & Klami, 2017). One could also consider employing synthetic gradients (Jaderberg et al., 2017) to limit the number of backpropagation steps during training. Krishnan et al. (2018) observe that it is more important to train with iterative refinement during earlier stages (we also observed this in preliminary experiments), and therefore annealing the number of refinement steps as training progresses could also speed up training.

Our approach is mainly applicable to variational families that avail themselves to differentiable optimization (e.g. gradient ascent) with respect to the ELBO, which include much recent work on employing more flexible variational families with VAEs. In contrast, VAE + SVI and VAE + SVI + KL are applicable to more general optimization algorithms.

6. Related Work

Our work is most closely related the line of work which uses a separate model to initialize variational parameters and subsequently updates them through an iterative procedure (Salakhutdinov & Larochelle, 2010; Cho et al., 2013; Salimans et al., 2015; Hjelm et al., 2016; Krishnan et al., 2018; Pu et al., 2017). Marino et al. (2018) utilize meta-learning to train an inference network which learns to perform iterative inference by training a deep model to output the variational parameters for each time step.

While differentiating through inference/optimization was initially explored by various researchers primarily outside the area of deep learning (Stoyanov et al., 2011; Domke, 2012; Brakel et al., 2013), they have more recently been explored in the context of hyperparameter optimization (Maclaurin et al., 2015) and as a differentiable layer of a deep model (Belanger et al., 2017; Kim et al., 2017; Metz et al., 2017; Amos & Kolter, 2017).

Initial work on VAE-based approaches to image modeling focused on simple generative models that assumed independence among pixels conditioned on the latent variable (Kingma & Welling, 2014; Rezende et al., 2014). More recent works have obtained substantial improvements in log-likelihood and sample quality through utilizing powerful autoregressive models (PixelCNN) as the generative model (Chen et al., 2017; Gulrajani et al., 2017).

In contrast, modeling text with VAEs has remained challenging. Bowman et al. (2016) found that using an LSTM generative model resulted in a degenerate case whereby the variational posterior collapsed to the prior and the generative model ignored the latent code (even with richer variational families). Many works on VAEs for text have thus made simplifying conditional independence assumptions (Miao et al., 2016; 2017), used less powerful generative models such as convolutional networks (Yang et al., 2017; Semeniuta et al., 2017), or combined a recurrent generative model with a topic model (Dieng et al., 2017; Wang et al., 2018). Note that unlike to sequential VAEs that employ different latent variables at each time step (Chung et al., 2015; Fraccaro et al., 2016; Krishnan et al., 2017; Serban et al., 2017; Goyal et al., 2017a), in this work we focus on modeling the entire sequence with a global latent variable.

Finally, since our work only addresses the *amortization gap* (the gap between the log-likelihood and the ELBO due to amortization) and not the *approximation gap* (due to the choice of a particular variational family) (Cremer et al., 2018), it can be combined with existing work on employing richer posterior/prior distributions within the VAE framework (Rezende & Mohamed, 2015; Kingma et al., 2016; Johnson et al., 2016; Tran et al., 2016; Goyal et al., 2017b; Guu et al., 2017; Tomczak & Welling, 2018).

7. Conclusion

This work outlines semi-amortized variational autoencoders, which combine amortized inference with local iterative refinement to train deep generative models of text and images. With the approach we find that we are able to train deep latent variable models of text with an expressive autoregressive generative model that does not ignore the latent code.

From the perspective of learning latent representations, one might question the prudence of using an autoregressive model that fully conditions on its entire history (as opposed to assuming some conditional independence) given that $p(\mathbf{x})$ can always be factorized as $\prod_{t=1}^T p(\mathbf{x}_t | \mathbf{x}_{<t})$, and therefore the model is *non-identifiable* (i.e. it does not have to utilize the latent variable). However in finite data regimes we might still expect a model that makes use of its latent variable to generalize better due to potentially better inductive bias (from the latent variable). Training generative models that both model the underlying data well and learn good latent representations is an important avenue for future work.

Acknowledgements

We thank Rahul Krishnan, Rachit Singh, and Justin Chiu for insightful comments/discussion. We additionally thank Zichao Yang for providing the text dataset. YK and AM are supported by Samsung Research. SW is supported by an Amazon AWS ML Award.

References

- Amos, Brandon and Kolter, J. Zico. OptNet: Differentiable Optimization as a Layer in Neural Networks. In *Proceedings of ICML*, 2017.
- Andrychowicz, Marcin, Denil, Misha, Gomez, Sergio, Hoffman, Matthew, Pfau, David, Schaul, Tom, and de Freitas, Nando. Learning to Learn by Gradient Descent by Gradient Descent. In *Proceedings of NIPS*, 2016.
- Belanger, David, Yang, Bishan, and McCallum, Andrew. End-to-end Learning for Structured Prediction Energy Networks. In *Proceedings of ICML*, 2017.
- Blei, David M., Kucukelbir, Alp, and McAuliffe, Jon D. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Bowman, Samuel R., Vilnis, Luke, Vinyal, Oriol, Dai, Andrew M., Jozefowicz, Rafal, and Bengio, Samy. Generating Sentences from a Continuous Space. In *Proceedings of CoNLL*, 2016.
- Brakel, Philemon, Stroobandt, Dirk, and Schrauwen, Benjamin. Training Energy-Based Models for Time-Series Imputation. *Journal of Machine Learning Research*, 14:2771–2797, 2013.
- Burda, Yuri, Grosse, Roger, and Salakhutdinov, Ruslan. Importance Weighted Autoencoders. In *Proceedings of ICLR*, 2015a.
- Burda, Yuri, Grosse, Roger, and Salakhutdinov, Ruslan. Accurate and Conservative Estimates of MRF Log-likelihood using Reverse Annealing. In *Proceedings of AISTATS*, 2015b.
- Chen, Xi, Kingma, Diederik P., Salimans, Tim, Duan, Yan, Dhariwal, Prafulla, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. Variational Lossy Autoencoder. In *Proceedings of ICLR*, 2017.
- Cho, Kyunghyun, Raiko, Tapani, Ilin, Alexander, and Karhunen, Juha. A Two-Stage Pretraining Algorithm for Deep Boltzmann Machines. In *Proceedings of ICANN*, 2013.
- Chung, Junyoung, Kastner, Kyle, Dinh, Laurent, Goel, Kratarth, Courville, Aaron, and Bengio, Yoshua. A Recurrent Latent Variable Model for Sequential Data. In *Proceedings of NIPS*, 2015.
- Cremer, Chris, Li, Xuechen, and Duvenaud, David. Inference Suboptimality in Variational Autoencoders. In *Proceedings of ICML*, 2018.
- Dieng, Adji B., Wang, Chong, Gao, Jianfeng, , and Paisley, John. TopicRNN: A Recurrent Neural Network With Long-Range Semantic Dependency. In *Proceedings of ICLR*, 2017.
- Domke, Justin. Generic Methods for Optimization-based Modeling. In *Proceedings of AISTATS*, 2012.
- Fraccaro, Marco, Sonderby, Soren Kaae, Paquet, Ulrich, and Winther, Ole. Sequential Neural Models with Stochastic Layers. In *Proceedings of NIPS*, 2016.
- Ghahramani, Zoubin and Beal, Matthew. Propagation algorithms for variational bayesian learning. In *Proceedings of NIPS*, 2001.
- Goyal, Anirudh, Sordoni, Alessandro, Cote, Marc-Alexandre, Ke, Nan Rosemary, and Bengio, Yoshua. Z-Forcing: Training Stochastic Recurrent Networks. In *Proceedings of NIPS*, 2017a.
- Goyal, Prasoon, Hu, Zhiting, Liang, Xiaodan, Wang, Chenyu, and Xing, Eric. Nonparametric Variational Auto-encoders for Hierarchical Representation Learning. In *Proceedings of ICCV*, 2017b.
- Gregor, Karol, Danihelka, Ivo, Graves, Alex, Rezende, Danilo Jimenez, and Wierstra, Daan. DRAW: A Recurrent Neural Network for Image Generation. In *Proceedings of ICML*, 2015.
- Gregor, Karol, Besse, Frederic, Rezende, Danilo Jimenez, Danihelka, Ivo, and Wierstra, Daan. Towards Conceptual Compression. In *Proceedings of NIPS*, 2016.
- Gulrajani, Ishaan, Kumar, Kundan, Ahmed, Faruk, Taiga, Adrien Ali, Visin, Francesco, Vazquez, David, and Courville, Aaron. PixelVAE: A Latent Variable Model for Natural Images. In *Proceedings of ICLR*, 2017.
- Guu, Kelvin, Hashimoto, Tatsunori B., Oren, Yonatan, and Liang, Percy. Generating Sentences by Editing Prototypes. *arXiv:1709.08878*, 2017.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of CVPR*, 2016.
- Hjelm, R Devon, Cho, Kyunghyun, Chung, Junyoung, Salakhutdinov, Russ, Calhoun, Vince, and Jojic, Nebojsa. Iterative Refinement of the Approximate Posterior for Directed Belief Networks. In *Proceedings of NIPS*, 2016.
- Hoffman, Matthew D., Blei, David M., Wang, Chong, and Paisley, John. Stochastic Variational Inference. *Journal of Machine Learning Research*, 13:1303–1347, 2013.
- Hornik, Kur, Stinchcombe, Maxwell, and White, Halber. Multilayer Feedforward Networks are Universal Approximators. *Neural Networks*, 2:359–366, 1989.
- Hu, Zhiting, Yang, Zichao, Liang, Xiaodan, Salakhutdinov, Ruslan, and Xing, Eric P. Toward Controlled Generation of Text. In *Proceedings of ICML*, 2017.
- Jaderberg, Max, Czarnecki, Wojciech Marian, Osindero, Simon, Vinyals, Oriol, Graves, Alex, Silver, David, and Kavukcuoglu, Koray. Decoupled Neural Interfaces using Synthetic Gradients. In *Proceedings of ICML*, 2017.
- Johnson, Matthew, Duvenaud, David K., Wiltchko, Alex, Adams, Ryan P., and Datta, Sandeep R. Composing Graphical Models with Neural Networks for Structured Representations and Fast Inference. In *Proceedings of NIPS*, 2016.
- Jordan, Michael, Ghahramani, Zoubin, Jaakkola, Tommi, and Saul, Lawrence. Introduction to Variational Methods for Graphical Models. *Machine Learning*, 37:183–233, 1999.
- Kim, Yoon, Denton, Carl, Hoang, Luong, and Rush, Alexander M. Structured Attention Networks. In *Proceedings of ICLR*, 2017.

- Kingma, Diederik P. and Welling, Max. Auto-Encoding Variational Bayes. In *Proceedings of ICLR*, 2014.
- Kingma, Diederik P., Salimans, Tim, and Welling, Max. Improving Variational Inference with Autoregressive Flow. In *Proceedings of ICLR Workshop*, 2016.
- Krishnan, Rahul G., Shalit, Uri, and Sontag, David. Structured Inference Networks for Nonlinear State Space Models. In *Proceedings of AAAI*, 2017.
- Krishnan, Rahul G., Liang, Dawen, and Hoffman, Matthew. On the Challenges of Learning with Inference Networks on Sparse, High-dimensional Data. In *Proceedings of AISTATS*, 2018.
- Lake, Brendan M., Salakhutdinov, Ruslan, and Tenenbaum, Joshua B. Human-level Concept Learning through Probabilistic Program Induction. *Science*, 350:1332–1338, 2015.
- LeCun, Yann, Simard, Patrice, and Pearlmutter, Barak. Automatic Learning Rate Maximization by On-line Estimation of the Hessians Eigenvectors. In *Proceedings of NIPS*, 1993.
- Maclaurin, Dougal, Duvenaud, David, and Adams, Ryan P. Gradient-based Hyperparameter Optimization through Reversible Learning. In *Proceedings of ICML*, 2015.
- Marino, Joseph, Yue, Yisong, and Mandt, Stephan. Iterative Amortized Inference. In *Proceedings of ICML*, 2018.
- Metz, Luke, Poole, Ben, Pfau, David, and Sohl-Dickstein, Jascha. Unrolled Generative Adversarial Networks. In *Proceedings of ICLR*, 2017.
- Miao, Yishu, Yu, Lei, and Blunsom, Phil. Neural Variational Inference for Text Processing. In *Proceedings of ICML*, 2016.
- Miao, Yishu, Grefenstette, Edward, and Blunsom, Phil. Discovering Discrete Latent Topics with Neural Variational Inference. In *Proceedings of ICML*, 2017.
- Mnih, Andriy and Gregor, Karol. Neural Variational Inference and Learning in Belief Networks. In *Proceedings of ICML*, 2014.
- Mueller, Jonas, Gifford, David, and Jaakkola, Tommi. Sequence to Better Sequence: Continuous Revision of Combinatorial Structures. In *Proceedings of ICML*, 2017.
- Pearlmutter, Barak A. Fast exact multiplication by the hessian. *Neural Computation*, 6(1):147–160, 1994.
- Pu, Yunchen, Gan, Zhe, Henao, Ricardo, Li, Chunyuan, Han, Shaobo, and Carin, Lawrence. VAE Learning via Stein Variational Gradient Descent. In *Proceedings of NIPS*, 2017.
- Ranganath, Rajesh, Gerrish, Sean, and Blei, David M. Black Box Variational Inference. In *Proceedings of AISTATS*, 2014.
- Ranganath, Rajesh, Tran, Dustin, and Blei, David M. Hierarchical Variational Models. In *Proceedings of ICML*, 2016.
- Rezende, Danilo J. and Mohamed, Shakir. Variational Inference with Normalizing Flows. In *Proceedings of ICML*, 2015.
- Rezende, Danilo Jimenez, Mohamed, Shakir, and Wierstra, Daan. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of ICML*, 2014.
- Rolfe, Jason Tyler. Discrete Variational Autoencoders. In *Proceedings of ICLR*, 2017.
- Sakaya, Joseph and Klami, Arto. Importance Sampled Stochastic Optimization for Variational Inference. In *Proceedings of UAI*, 2017.
- Salakhutdinov, Ruslan and Larochelle, Hugo. Efficient Learning of Deep Boltzmann Machines. In *Proceedings of AISTATS*, 2010.
- Salimans, Tim, Kingma, Diederik, and Welling, Max. Markov Chain Monte Carlo and Variational Inference: Bridging the Gap. In *Proceedings of ICML*, 2015.
- Schmidt, Mark, Roux, Nicolas Le, and Bach, Francis. Minimizing Finite Sums with the Stochastic Average Gradient. *arXiv:1309.2388*, 2013.
- Semeniuta, Stanislau, Severyn, Aliaksei, and Barth, Erhardt. A Hybrid Convolutional Variational Autoencoder for Text Generation. In *Proceedings of EMNLP*, 2017.
- Serban, Iulian Vlad, Sordani, Alessandro, Ryan Lowe, Laurent Charlin, Pineau, Joelle, Courville, Aaron, and Bengio, Yoshua. A Hierarchical Latent Variable Encoder-Decoder Model for Generating Dialogues. In *Proceedings of AAAI*, 2017.
- Sønderby, Casper Kaae, Raiko, Tapani, Maaløe, Lars, Sønderby, Søren Kaae, and Winther, Ole. Ladder Variational Autoencoders. In *Proceedings of NIPS*, 2016.
- Stoyanov, Veselin, Ropson, Alexander, and Eisner, Jason. Empirical Risk Minimization of Graphical Model Parameters Given Approximate Inference, Decoding, and Model Structure. In *Proceedings of AISTATS*, 2011.
- Tomczak, Jakub M. and Welling, Max. VAE with a VampPrior. In *Proceedings of AISTATS*, 2018.
- Tran, Dustin, Ranganath, Rajesh, and Blei, David M. The Variational Gaussian Process. In *Proceedings of ICLR*, 2016.
- van den Oord, Aaron, Kalchbrenner, Nal, Vinyals, Oriol, Espenholt, Lasse, Graves, Alex, and Kavukcuoglu, Koray. Conditional Image Generation with PixelCNN Decoders. In *Proceedings of NIPS*, 2016.
- Wainwright, Martin J. and Jordan, Michael I. Introduction to Variational Methods for Graphical Models. *Foundations and Trends in Machine Learning*, 1:1–103, 2008.
- Wang, Wenlin, Gan, Zhe, Wang, Wenqi, Shen, Dinghan, Huang, Jiaji, Ping, Wei, Satheesh, Sanjeev, and Carin, Lawrence. Topic Compositional Neural Language Model. In *Proceedings of AISTATS*, 2018.
- Yang, Zichao, Hu, Zhiting, Salakhutdinov, Ruslan, and Berg-Kirkpatrick, Taylor. Improved Variational Autoencoders for Text Modeling using Dilated Convolutions. In *Proceedings of ICML*, 2017.
- Zhang, Cheng, Butepage, Judith, Kjellstrom, Hedvig, and Mandt, Stephan. Advances in Variational Inference. *arXiv:1711.05597*, 2017.
- Zhao, Shengjia, Song, Jiaming, and Ermon, Stefano. Towards Deeper Understanding of Variational Autoencoding Models. In *Proceedings of ICML*, 2017.

Supplementary Materials for Semi-Amortized Variational Autoencoders

A. Training Semi-Amortized Variational Autoencoders with Gradient Clipping

For stable training we found it crucial to modify Algorithm 1 to clip the gradients at various stages. This is shown in Algorithm 2, where we have a clipping parameter η . The $\text{clip}(\cdot)$ function is given by

$$\text{clip}(\mathbf{u}, \eta) = \begin{cases} \frac{\eta}{\|\mathbf{u}\|_2} \mathbf{u}, & \text{if } \|\mathbf{u}\|_2 > \eta \\ \mathbf{u}, & \text{otherwise} \end{cases}$$

We use $\eta = 5$ in all experiments. The finite difference estimation itself also uses gradient clipping. See <https://github.com/harvardnlp/sa-vae/blob/master/optim.n2n.py> for the exact implementation.

B. Experimental Details

For all the variational models we use a spherical Gaussian prior. The variational family is the diagonal Gaussian parameterized by the vector of means and log variances. For models trained with SVI the initial variational parameters are randomly initialized from a Gaussian with standard deviation equal to 0.1.

B.1. Synthetic Data

We generate synthetic data points according to the following generative process:

$$z_1, z_2 \sim \mathcal{N}(0, 1) \quad \mathbf{h}_t = \text{LSTM}(\mathbf{h}_{t-1}, \mathbf{x}_t) \\ \mathbf{x}_{t+1} \sim \text{softmax}(\text{MLP}([\mathbf{h}_t, z_1, z_2]))$$

Here LSTM is a one-layer LSTM with 100 hidden units where the input embedding is also 100-dimensional. The initial hidden/cell states are set to zero, and we generate for 5 time steps for each example (so $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_5]$). The MLP consists of a single affine transformation to project out to the vocabulary space, which has 1000 tokens. LSTM/MLP parameters are randomly initialized with $\mathcal{U}(-1, 1)$, except for the part of the MLP that directly connects to the latent variables, which is initialized with $\mathcal{U}(-5, 5)$. This is done to make sure that the latent variables have more influence in predicting \mathbf{x} . We generate 5000 training/validation/test examples.

When we learn the generative model the LSTM is initialized over $\mathcal{U}(-0.1, 0.1)$. The inference network is also a one-layer LSTM with 100-dimensional hidden units/input

Algorithm 2 Semi-Amortized Variational Autoencoders with Gradient Clipping

Input: inference network ϕ , generative model θ , inference steps K , learning rate α , momentum γ , loss function $f(\lambda, \theta, \mathbf{x}) = -\text{ELBO}(\lambda, \theta, \mathbf{x})$, gradient clipping parameter η

Sample $\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})$

$\lambda_0 \leftarrow \text{enc}(\mathbf{x}; \phi)$

$v_0 \leftarrow 0$

for $k = 0$ **to** $K - 1$ **do**

$v_{k+1} \leftarrow \gamma v_k - \text{clip}(\nabla_{\lambda} f(\lambda_k, \theta, \mathbf{x}), \eta)$

$\lambda_{k+1} \leftarrow \lambda_k + \alpha v_{k+1}$

end for

$\mathcal{L} \leftarrow f(\lambda_K, \theta, \mathbf{x})$

$\bar{\lambda}_K \leftarrow \nabla_{\lambda} f(\lambda_K, \theta, \mathbf{x})$

$\bar{\theta} \leftarrow \nabla_{\theta} f(\lambda_K, \theta, \mathbf{x})$

$\bar{v}_K \leftarrow 0$

for $k = K - 1$ **to** 0 **do**

$\bar{v}_{k+1} \leftarrow \bar{v}_{k+1} + \alpha \bar{\lambda}_{k+1}$

$\bar{\lambda}_k \leftarrow \bar{\lambda}_{k+1} - \text{H}_{\lambda, \lambda} f(\lambda_k, \theta, \mathbf{x}) \bar{v}_{k+1}$

$\bar{\lambda}_k \leftarrow \text{clip}(\bar{\lambda}_k, \eta)$

$\bar{\theta} \leftarrow \bar{\theta} - \text{clip}(\text{H}_{\theta, \lambda} f(\lambda_k, \theta, \mathbf{x}) \bar{v}_{k+1}, \eta)$

$\bar{v}_k \leftarrow \gamma \bar{v}_{k+1}$

end for

$\frac{d\mathcal{L}}{d\theta} \leftarrow \bar{\theta}$

$\frac{d\mathcal{L}}{d\phi} \leftarrow \frac{d\lambda_0}{d\phi} \bar{\lambda}_0$

Update θ, ϕ based on $\frac{d\mathcal{L}}{d\theta}, \frac{d\mathcal{L}}{d\phi}$

embeddings, where the variational parameters are predicted via an affine transformation on the final hidden state of the encoder. All models are trained with stochastic gradient descent with batch size 50, learning rate 1.0, and gradient clipping at 5. The learning rate starts decaying by a factor of 2 each epoch after the first epoch at which validation performance does not improve. This learning rate decay is not triggered for the first 5 epochs. We train for 20 epochs, which was enough for convergence of all models. For SVI/SA-VAE we perform 20 steps of iterative inference with stochastic gradient descent and learning rate 1.0 with gradient clipping at 5.

B.2. Text

We use the same model architecture as was used in Yang et al. (2017). The inference network and the generative model are both one-layer LSTMs with 1024-dimensional

INFERENCE NETWORK	3-LAYER RESNET		2-LAYER MLP	
MODEL	VAE	SA-VAE	VAE	SA-VAE
DATA SIZE: 25%	92.21 (0.81)	91.89 (2.31)	92.33 (0.27)	92.03 (1.32)
DATA SIZE: 50%	91.38 (0.77)	91.01 (2.54)	91.40 (0.51)	91.10 (1.48)
DATA SIZE: 75%	90.82 (1.06)	90.51 (2.07)	90.90 (0.45)	90.67 (1.34)
DATA SIZE: 100%	90.43 (0.98)	90.05 (2.78)	90.56 (0.61)	90.25 (1.77)
1-LAYER PIXELCNN	96.53 (10.36)	96.01 (10.93)	98.30 (8.87)	96.43 (10.14)
3-LAYER PIXELCNN	93.75 (7.10)	93.16 (8.73)	94.45 (5.46)	93.55 (7.20)
6-LAYER PIXELCNN	91.24 (3.25)	90.79 (4.44)	91.40 (2.06)	91.01 (3.27)
9-LAYER PIXELCNN	90.54 (1.78)	90.28 (3.02)	90.72 (1.14)	90.34 (2.26)
12-LAYER PIXELCNN	90.43 (0.98)	90.05 (2.78)	90.56 (0.61)	90.25 (1.77)

Table 4. Upper bounds on negative log-likelihood (i.e. negative ELBO) of VAE/SA-VAE trained on OMNIGLOT, where we vary the capacity of the inference network (3-layer ResNet vs 2-layer MLP). KL portion of the loss is shown in parentheses. (Top) Here we vary the training set size from 25% to 100%, and use a 12-layer PixelCNN as the generative model. (Bottom) Here we fix the training set size to be 100%, and vary the capacity of the generative model.

hidden states where the input word embedding is 512-dimensional. We use the final hidden state of the encoder to predict (via an affine transformation) the vector of variational means and log variances. The latent space is 32-dimensional. The sample from the variational posterior is used to initialize the initial hidden state of the generative LSTM (but not the cell state) via an affine transformation, and additionally fed as input (i.e. concatenated with the word embedding) at each time step. There are dropout layers with probability 0.5 between the input-to-hidden layer and the hidden-to-output layer on the generative LSTM only.

The data contains 100000/10000/10000 train/validation/test examples with 20000 words in the vocabulary. All models are trained with stochastic gradient descent with batch size 32 and learning rate 1.0, where the learning rate starts decaying by a factor of 2 each epoch after the first epoch at which validation performance does not improve. This learning rate decay is not triggered for the first 15 epochs to ensure adequate training. We train for 30 epochs or until the learning rate has decayed 5 times, which was enough for convergence for all models. Model parameters are initialized over $\mathcal{U}(-0.1, 0.1)$ and gradients are clipped at 5. We employ a KL-cost annealing schedule whereby the multiplier on the KL-cost term is increased linearly from 0.1 to 1.0 each batch over 10 epochs. For models trained with iterative inference we perform SVI via stochastic gradient descent with momentum 0.5 and learning rate 1.0. Gradients are clipped after each step of SVI (also at 5).

B.3. Images

The preprocessed OMNIGLOT dataset does not have a standard validation split so we randomly pick 2000 images from training as validation. As with previous works the pixel value is scaled to be between 0 and 1 and interpreted as probabilities, and the images are dynamically binarized dur-

ing training.

Our inference network consists of 3 residual blocks where each block is made up of a standard residual layer (i.e. two convolutional layers with 3×3 filters, ReLU nonlinearities, batch normalization, and residual connections) followed by a downsampling convolutional layer with filter size and stride equal to 2. These layers have 64 feature maps. The output of residual network is flattened and then used to obtain the variational means/log variances via an affine transformation.

The sample from the variational distribution (which is 32-dimensional) is first projected out to the image spatial resolution with 4 feature maps (i.e. $4 \times 28 \times 28$) via a linear transformation, then concatenated with the original image, and finally fed as input to a 12-layer Gated PixelCNN (van den Oord et al., 2016). The PixelCNN has three 9×9 layers, followed by three 7×7 layers, then three 5×5 layers, and finally three 3×3 layers. All the layers have 32 feature maps, and there is a final 1×1 convolutional layer followed by a sigmoid nonlinearity to produce a distribution over binary output. The layers are appropriately masked to ensure that the distribution over each pixel is conditioned only on the pixels left/top of it. We train with Adam with learning rate 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$ for 100 epochs with batch size of 50. Gradients are clipped at 5.

For models trained with iterative inference we perform SVI via stochastic gradient descent with momentum 0.5 and learning rate 1.0, with gradient clipping (also at 5).

C. Data Size/Model Capacity

In Table 4 we investigate the performance of VAE/SA-VAE as we vary the capacity of the inference network, size of the training set, and the capacity of the generative model. The MLP inference network has two ReLU layers with 128 hidden units. For varying the PixelCNN generative model,

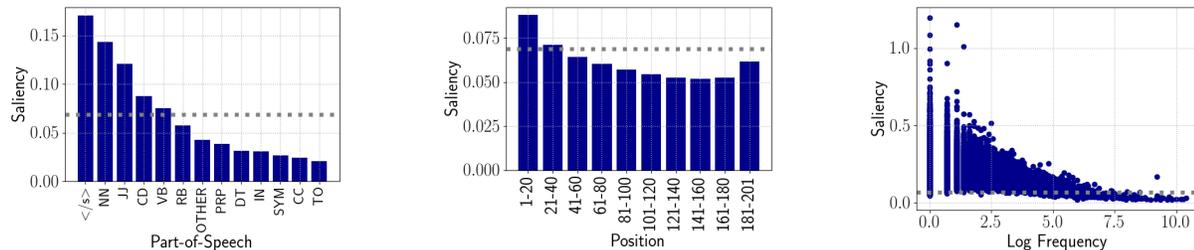


Figure 5. Input saliency by part-of-speech tag (left), position (center), and log frequency (right). The dotted gray line in each plot shows the average saliency across all words.

we sequentially remove layers from our baseline 12-layer model starting from the bottom (so the 9-layer PixelCNN has three 7×7 layers, three 5×5 layers, three 3×3 layers, all with 32 feature maps).

Intuitively, we expect iterative inference to help more when the inference network and the generative model are less powerful, and we indeed see this in Table 4. Further, one might expect SA-VAE to be more helpful in small-data regimes as it is harder for the inference network amortize inference and generalize well to unseen data. However we find that SA-VAE outperforms VAE by a similar margin across all training set sizes.

Finally, we observe that across all scenarios the KL portion of the loss is much higher for models trained with SA-VAE, indicating that these models are learning generative models that make more use of the latent representations.

D. Input Saliency Analysis

In Figure 5 we show the input saliency by part-of-speech tag (left), position (center), and frequency (right). Input saliency of a token \mathbf{x}_t is defined as:

$$\left\| \mathbb{E}_{q(\mathbf{z}; \lambda)} \left[\frac{d \|\mathbf{z}\|_2}{d \mathbf{w}_t} \right] \right\|_2$$

Here \mathbf{w}_t is the encoder word embedding for \mathbf{x}_t . Part-of-speech tagging is done using NLTK.