

Boltzmann Encoded Adversarial Machines

Charles K. Fisher

Unlearn.AI, Inc.
San Francisco, CA 94108, USA

DRCKF@UNLEARN.AI

Aaron M. Smith

Unlearn.AI, Inc.
San Francisco, CA 94108, USA

DRAMS@UNLEARN.AI

Jonathan R. Walsh

Unlearn.AI, Inc.
San Francisco, CA 94108, USA

DRJRW@UNLEARN.AI

Editor: N.A.

Abstract

Restricted Boltzmann Machines (RBMs) are a class of generative neural network that are typically trained to maximize a log-likelihood objective function. We argue that likelihood-based training strategies may fail because the objective does not sufficiently penalize models that place a high probability in regions where the training data distribution has low probability. To overcome this problem, we introduce Boltzmann Encoded Adversarial Machines (BEAMs). A BEAM is an RBM trained against an adversary that uses the hidden layer activations of the RBM to discriminate between the training data and the probability distribution generated by the model. We present experiments demonstrating that BEAMs outperform RBMs and GANs on multiple benchmarks.

Keywords: Generative Adversarial Networks, Restricted Boltzmann Machines

1. Introduction

A machine learning model is *generative* if it learns to draw new samples from an unknown probability distribution. Generative models have two important applications. First, generative models enable simulations of systems with unknown, or very complicated, mechanistic laws. For example, generative models can be used to design molecular compounds with desired properties Kadurin et al. (2017). Second, in the process of learning to generate samples from a distribution a generative model must learn a useful representation of the data. Therefore, generative models enable unsupervised learning with unlabeled data (Hinton and Sejnowski (1999)).

The last decade has produced revolutionary advances in machine learning, largely due to progress in training neural networks. Much of this progress has been on *discriminative* models rather than generative models. Still, neural generative models such as Restricted Boltzmann Machines (RBMs) (Hinton and Salakhutdinov (2006); Salakhutdinov and Hinton (2009)), Variational Autoencoders (VAEs) (Kingma and Welling (2013); Rolfe (2016); Kuleshov and Ermon (2017)), and Generative Adversarial Networks (GANs) (Goodfellow

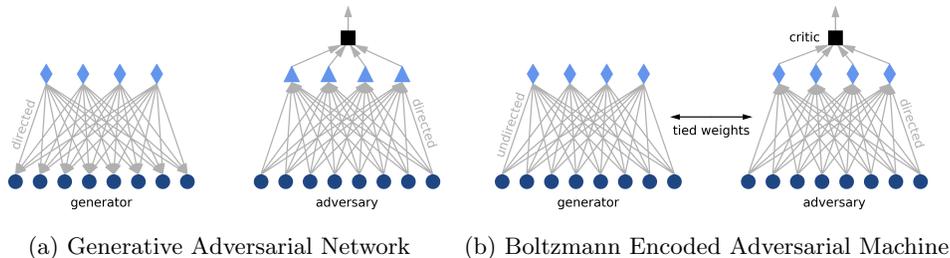


Figure 1: **Architecture of a BEAM.** (a) The generator of a GAN is a feed-forward neural network that transforms random noise into an image, and the adversary is a feed-forward neural network classifies the input image. (b) A BEAM uses an RBM generator trained to minimize an objective function that combines the negative log-likelihood and an adversarial loss. The adversarial loss is computed by a critic trained on the activations of the hidden units of the generator.

et al. (2014)) have demonstrated promising results on a number of problems. GANs, in particular, are generally regarded as the current state-of-the-art (Karras et al. (2017)).

Unlike most other generative models, GANs are trained to minimize a distance between the data and model distributions rather than to maximize the likelihood of the data under the model (Arjovsky and Bottou (2017); Nowozin et al. (2016)). As a result of the form of this distance function, and because they are built on feedforward neural networks, typical formulations of GANs can be trained using standard backpropagation (Rumelhart et al. (1986)). However, GANs have their drawbacks. GAN training can be difficult and unstable (Arjovsky and Bottou (2017); Arjovsky et al. (2017a)). Moreover, although one of the main advantages of GANs is that they can be trained end-to-end using backpropagation, recent state-of-the-art approaches have used a layerwise training strategy (Karras et al. (2017)) reminiscent of methods used to train Deep Boltzmann Machines (Hinton and Salakhutdinov (2012)).

The popularity of RBM-based generative models, including Deep Belief Networks and Deep Boltzmann Machines, has faded in recent years. The charge is that other approaches, especially GANs, simply work better in practice. However, RBM based architectures do have some advantages. For example, RBMs can be easily adapted for use on multimodal data sets (Srivastava and Salakhutdinov (2012)) and on time series (Taylor et al. (2007); Taylor and Hinton (2009); Sutskever et al. (2009)) without major modifications, and RBMs allow one to perform both generation and inference with a single model. Given that RBMs and derived models generally have sufficient representational power to learn essentially any distribution (Le Roux and Bengio (2008)), the difficulties must arise during training.

In this work, we take inspiration from GANs to propose a new method for training RBMs. We call the resulting model a Boltzmann Encoded Adversarial Machine (BEAM; see Figure 1). While the adversarial concept used in BEAMs is similar to GANs, there are some distinct features. The primary one is that the adversary operates on the hidden layer activations of the RBM. Because the latent variable representation from the RBM is a consolidated representation of the visible units, simple adversaries – even ones that do not

need to be trained – are often sufficient to obtain good results. This makes training simple and stable. Furthermore, we obtain our best results by optimizing a convex combination of the log-likelihood and adversarial objectives. The component of the objective from the log-likelihood allows the training data to play an active role in determining the gradient (while it only plays a passive role as part of the discriminator in the adversarial gradient, as it also does in GANs).

BEAMs achieve excellent results on a variety of applications, from low dimensional benchmark datasets to higher dimensional applications such as image generation, outperforming GANs of similar or higher complexity. These results indicate that BEAMs provide a powerful approach to unsupervised learning.

This paper is structured as follows. We begin with a brief review of RBMs, then discuss some problems with maximum likelihood training of RBMs and go on to define and describe BEAMs. Finally, we present the results of experiments comparing RBMs, GANs, and BEAMs and discuss.

2. Theory and Methods

2.1 Restricted Boltzmann Machines

An RBM is an energy based model with two layers of neurons. The visible units \mathbf{v} describe the data and the hidden units \mathbf{h} capture interactions between the visible units. The joint probability distribution $p(\mathbf{v}, \mathbf{h}) = Z^{-1} e^{-E(\mathbf{v}, \mathbf{h})}$ is defined by an energy function:

$$E(\mathbf{v}, \mathbf{h}) = - \sum_i a_i(v_i) - \sum_\mu b_\mu(h_\mu) - \sum_{i\mu} W_{i\mu} \frac{v_i}{\sigma_i^2} \frac{h_\mu}{\epsilon_\mu^2} \quad (1)$$

with a partition function $Z = \int d\mathbf{v} d\mathbf{h} \exp(-E(\mathbf{v}, \mathbf{h}))$. This formulation, where $a_i(\cdot)$ and $b_\mu(\cdot)$ are generic functions and σ_i and ϵ_μ are scale parameters, is a flexible way of writing a generic RBM that encompasses common models such as Bernoulli RBMs and Gaussian RBMs. The key feature of an RBM is the conditional independence of the layers, i.e. $p(\mathbf{v}|\mathbf{h}) = \prod_i p_i(v_i|\mathbf{h})$ and $p(\mathbf{h}|\mathbf{v}) = \prod_\mu p_\mu(h_\mu|\mathbf{v})$, which allows one to sample from the distribution using block Gibbs sampling.

RBMs are typically trained to maximize the log-likelihood $\mathcal{L} = \langle \log \int d\mathbf{h} p(\mathbf{v}, \mathbf{h}) \rangle_{data}$ using algorithms such as Persistent Contrastive Divergence (PCD) (Tieleman (2008); Hinton (2006)). The derivative of the log-likelihood with respect to a model parameter θ takes the form (Ackley et al. (1985)):

$$\partial_\theta \mathcal{L} = \langle -\partial_\theta E(\mathbf{v}, \mathbf{h}) \rangle_{data} - \langle -\partial_\theta E(\mathbf{v}, \mathbf{h}) \rangle_{model}. \quad (2)$$

The two averages are computed using samples from the data set and samples drawn from the model by Gibbs sampling, respectively. We refer the reader to foundational works such as Hinton (2010) for more detail.

2.2 The Problem with Maximum Likelihood

A generative model defined by parameters θ describes the probability of observing a visible state \mathbf{v} . Therefore, training a generative model involves minimizing a distance between

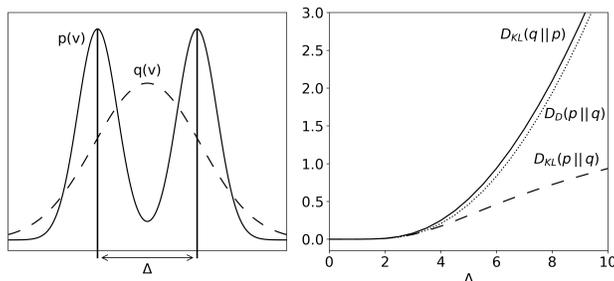


Figure 2: **Comparison of distances between distributions.** We consider the distance between $p(v)$, a mixture of two Gaussian distributions separated by a distance Δ , and $q(v)$, a single Gaussian distribution with the same mean and standard deviation as $p(v)$. The forward KL divergence $D_{\text{KL}}(p \parallel q)$ increases slowly as Δ increases, while the reverse KL divergence $D_{\text{KL}}(q \parallel p)$ and discriminator divergence $D_D(p \parallel q)$ increase rapidly.

the distribution of the data, $p_d(\mathbf{v})$, and the distribution defined by the model, $p_\theta(\mathbf{v})$. The traditional algorithms for training RBMs maximize the log-likelihood, which is equivalent to minimizing the forward Kullback-Liebler (KL) divergence (Kullback and Leibler (1951)):

$$D_{\text{KL}}(p_d \parallel p_\theta) = \int d\mathbf{v} p_d(\mathbf{v}) \log \left(\frac{p_d(\mathbf{v})}{p_\theta(\mathbf{v})} \right). \quad (3)$$

To illustrate some problems with maximum likelihood, we will compare the forward KL divergence to the reverse KL divergence,

$$D_{\text{KL}}(p_\theta \parallel p_d) = \int d\mathbf{v} p_\theta(\mathbf{v}) \log \left(\frac{p_\theta(\mathbf{v})}{p_d(\mathbf{v})} \right). \quad (4)$$

The forward KL divergence, $D_{\text{KL}}(p_d \parallel p_\theta)$, accumulates differences between the data and model distributions *weighted by the probability under the data distribution*. The reverse KL divergence, $D_{\text{KL}}(p_\theta \parallel p_d)$, accumulates differences between the data and model distributions *weighted by the probability under the model distribution*. As a result, the forward KL divergence strongly punishes models that *underestimate* the probability of the data, whereas the reverse KL divergence strongly punishes models that *overestimate* the probability of the data. Figure 2 illustrates the difference between the metrics on a simple mixture of two Gaussians.

There are a variety of sources of stochasticity that enter into the training of an RBM. Mandt et al. (2017) showed that stochastic gradient descent with a constant stepsize samples from an Ornstein-Uhlenbeck process that approximates the posterior distribution of the model parameters. Moreover, the moments used to compute the model gradients have to be estimated using random sampling by Markov Chain Monte Carlo methods. The stochasticity implies that different models may become statistically indistinguishable if the differences in their log-likelihoods are smaller than the errors in estimating them. This creates an *entropic force* because there will be many more models with a small $D_{\text{KL}}(p_d \parallel p_\theta)$ than there are models with both a small $D_{\text{KL}}(p_d \parallel p_\theta)$ and $D_{\text{KL}}(p_\theta \parallel p_d)$. As a result, training an RBM using a standard approach with PCD decreases $D_{\text{KL}}(p_d \parallel p_\theta)$ (as it should) *but*

tends to increase $D_{\text{KL}}(p_\theta \| p_d)$. This leads to distributions with spurious modes and/or to distributions that are oversmoothed.

2.3 Advantages of Adversarial Training

One can imagine overcoming the limitations of maximum likelihood training of RBMs by minimizing a combination of the forward and reverse KL divergences. Unfortunately, computing the reverse KL divergence requires knowledge of p_d , which is unknown. Therefore, we introduce a new type of f-divergence that we call a *discriminator divergence*

$$D_D(p_d \| p_\theta) := - \int d\mathbf{v} p_\theta(\mathbf{v}) \log \left(\frac{2p_d(\mathbf{v})}{p_d(\mathbf{v}) + p_\theta(\mathbf{v})} \right), \quad (5)$$

Notice that the optimal discriminator between p_d and p_θ will assign a posterior probability

$$p(\text{data}|\mathbf{v}) = \frac{p_d(\mathbf{v})}{p_d(\mathbf{v}) + p_\theta(\mathbf{v})} \quad (6)$$

that the sample \mathbf{v} was drawn from the data distribution. Therefore, we can write the discriminator divergence as

$$D_D(p_d \| p_\theta) = -\log 2 - \int d\mathbf{v} p_\theta(\mathbf{v}) \log(p(\text{data}|\mathbf{v})) \quad (7)$$

to show that it measures the probability that the optimal discriminator will incorrectly classify a sample drawn from the model distribution as coming from the data distribution.

The discriminator divergence belongs to the class of f-divergences defined as $D_f(p||q) := \int dx q(x) f(p(x)/q(x))$. The function that defines the discriminator divergence is

$$f(t) = \log \left(\frac{t+1}{2t} \right) \quad (8)$$

which is convex with $f(1) = 0$ as required. It is easy to show that the discriminator divergence upper bounds the reverse KL divergence:

$$\begin{aligned} \log 2 + D_D(p_d \| p_\theta) &= \int d\mathbf{v} p_\theta(\mathbf{v}) \log \left(1 + \frac{p_\theta(\mathbf{v})}{p_d(\mathbf{v})} \right) \\ &\geq D_{\text{KL}}(p_\theta \| p_d). \end{aligned}$$

We introduce this relationship because we usually do not have access to $p_d(\mathbf{v})$ directly and cannot compute the reverse KL divergence. However, we can train a discriminator to approximate Equation 6 and, therefore, can approximate the discriminator divergence.

A generator that is able to trick the discriminator so that $p(\text{data}|\mathbf{v}) \approx 1$ for all samples drawn from p_θ will have a low discriminator divergence. The discriminator divergence closely mirrors the reverse KL divergence and strongly punishes models that *overestimate* the probability of the data (Figure 2). Therefore, as with GANs, we hypothesized that it may be possible to improve the training of RBMs using an adversary. Some previous research in this direction includes the Wasserstein RBM (Montavon et al. (2016)) and Associate Adversarial Networks (Arıcı and Celikyilmaz (2016)).

2.4 Boltzmann Encoded Adversarial Machines (BEAMs)

We introduce a method – called a Boltzmann Encoded Adversarial Machine (BEAM) – for training an RBM against an adversary. A BEAM minimizes a loss function that is a combination of the negative log-likelihood and an adversarial loss. The adversarial component ensures that BEAM training performs a simultaneous minimization of both the forward and reverse KL divergences, which prevents the oversmoothing problem observed with regular RBMs.

The architecture of a BEAM is very simple, and is illustrated in Figure 1. The RBM (the generative model) is trained with an objective,

$$\mathcal{C} = -\gamma\mathcal{L} - (1 - \gamma)\mathcal{A}, \quad (9)$$

that includes a contribution from an adversarial term, \mathcal{A} . In theory, the adversary could be any model that can be trained to approximate the optimal discriminator.

We take inspiration from GANs and train the RBM against a critic function. However, we use a critic function $T(\mathbf{h})$ that acts on the hidden unit activations rather than the visible units. That is, the adversary uses same architecture and weights as the RBM, and encodes visible units into hidden unit activations. These hidden unit activations, computed for both the data and fantasy particles sampled from the RBM, are used by a critic to estimate the distance between the data and model distributions. Thus, the BEAM adversarial term is

$$\mathcal{A} := \int d\mathbf{h} p_{\theta}(\mathbf{h}) T(\mathbf{h}). \quad (10)$$

This term has a straightforward interpretation: for any sensible critic, it is minimizing the distance between the marginal distributions of the hidden units under the data and model distributions.

Maximizing the accuracy of the discriminator between the data and model distributions (on the hidden units) would lead to the optimal discriminator:

$$p(\text{data}|\mathbf{h}) = \frac{p_d(\mathbf{h})}{p_d(\mathbf{h}) + p_{\theta}(\mathbf{h})} \quad (11)$$

where $p_d(\mathbf{h}) := \int d\mathbf{v} p_{\theta}(\mathbf{h}|\mathbf{v}) p_d(\mathbf{v})$. Therefore, we could directly minimize the discriminator divergence by defining the critic as $\tilde{T}(\mathbf{h}) = \log 2 + \log(p(\text{data}|\mathbf{h}))$. In practice, however, we found that we obtain better results with a linear critic:

$$T(\mathbf{h}) = 2p(\text{data}|\mathbf{h}) - 1. \quad (12)$$

Therefore, all experiments that follow use a linear critic. We use the $2p - 1$ form so that the sign of the critic indicates the best guess of the optimal discriminator, but this choice is not important since it only ends up scaling the derivative by a factor of two.

In practice, of course, we don't have access to the optimal discriminator. The usual remedy for GANs is to co-train a neural network to approximate it. That is, the min-max game between the generator and discriminator can be recast as minimizing a bound on a particular f-divergence (Nowozin et al. (2016)). In our case, we hypothesized that a simple approximation to the optimal discriminator will be sufficient because we are working with the

hidden unit activities of the RBM generator rather than the visible units. Therefore, we simply approximate the optimal critic using nearest neighbor methods. In our examples, we simply store the data and fantasy particles from the previous minibatch and use a distance-weighted nearest neighbor approximation.

A BEAM can be trained using stochastic gradient descent by computing model averages from persistent fantasy particles in the same way as with maximum likelihood training of an RBM. The derivative of the adversarial term with respect to a model parameter θ is

$$\partial_{\theta}\mathcal{A} = \text{Cov}_{\theta}[T(\mathbf{h}), -\partial_{\theta}E(\mathbf{v}, \mathbf{h})], \quad (13)$$

where the covariance is computed with respect to the model distribution p_{θ} . A derivation of this result is presented in the Supplementary Material. It is also possible to define a critic on the visible units directly, or to use some other method other than a nearest neighbor approximation. We present some comparisons of BEAMs with other critics in Appendix 4.

In the context of most formulations of GANs – which use feed-forward neural networks for both the generator and the discriminator – one could say that BEAMs use the RBM as both the generator and as a feature extractor for the adversary. This double-usage allows us to reuse a single set of fantasy particles for multiple steps of the training algorithm. Specifically, we maintain a single set of M persistent fantasy particles that are updated k times per gradient evaluation. The same set of fantasy particles are used to compute the log-likelihood derivative (Equation 2) and the adversarial derivative (Equation 13). Then, these fantasy particles replace the fantasy particles from the previous gradient evaluation in the nearest neighbor estimates of the critic value. Reusing the fantasy particles for each step means that BEAM training has roughly the same computational cost as training an RBM with PCD.

2.5 Nearest Neighbor Critics

Suppose $X = \{x_1, \dots, x_N\}$ are i.i.d. samples from an unknown probability distribution with density function $p(x)$ in \mathbb{R}^n . One simple way to estimate $p(x)$ at an arbitrary point x is to make use of a k -nearest-neighbor estimate. Specifically: fix some positive integer k and compute the k nearest neighbors to x in X . Define d_k to be the distance between x and the furthest of the nearest-neighbors. Then estimate the density $p(x)$ to be the density of the uniform distribution on a ball of radius d_k . That is,

$$p(x) = k \left(\frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2} + 1)} d_k^n \right)^{-1}. \quad (14)$$

Now denote by $p_{\theta}(x)$ and $p_d(x)$ the unknown densities of the model and data distributions respectively. Suppose $X = \{x_1, \dots, x_{2N}\}$ is a collection of i.i.d. samples exactly half of which are drawn from p_{θ} and half from p_d . We can use the same idea to estimate the ratio $\frac{p_d}{p_d+p_{\theta}}(x)$. Fix some k and compute the k nearest neighbors in X , denoting by d_k the distance to the furthest. Then we estimate the denominator as in (14). Let j be the number of nearest neighbors that come from p_d as opposed to p_{θ} . The numerator then can be estimated as uniform on the same size ball with only j/k of the density of the denominator. As a result the desired estimate is simply the ratio j/k .

We put this concept in action by defining the *nearest-neighbor critic*. Suppose that we have cached a minibatch of samples from the model and a minibatch of samples from the training dataset. For any new sample x we can compute the k -nearest neighbors from the joined minibatches for some fixed k – we generally use $k = 5$ in examples. Then the nearest-neighbor critic is defined as the function that assigns to x the ratio j/k , where j is the number of nearest neighbors originating from the data minibatch as opposed to the model minibatch.

$$T_{NN}(x) := 2\frac{j}{k} - 1. \quad (15)$$

The *distance-weighted nearest-neighbor critic* is a generalization that attempts to add some continuity to the nearest-neighbor critic by applying an inverse distance weighting to the ratio count. Specifically, let $\{d_0, \dots, d_k\}$ be the distances of the k -nearest neighbors in X to some x , with $\{d_0, \dots, d_j\}$ the distances for the neighbors originating from the data samples and $\{d_{j+1}, \dots, d_k\}$ the distances for the neighbors originating from the model samples. Then distance-weighted nearest-neighbor critic is defined as:

$$T_{DNN}(x) := 2\frac{\sum_{i=1}^j \frac{1}{d_i + \epsilon}}{\sum_{i=1}^k \frac{1}{d_i + \epsilon}} - 1, \quad (16)$$

where ϵ regularizes the inverse distance.

2.6 Temperature Driven Sampling

Finally, we use a simple trick to improve the mixing of the RBM while sampling the fantasy particles. We assign each fantasy particle an independently sampled inverse temperature β and define the probability as $p(\mathbf{v}, \mathbf{h}) = Z^{-1}e^{-\beta E(\mathbf{v}, \mathbf{h})}$. The inverse temperature is drawn from an autoregressive Gamma process (Gouriéroux and Jasiak (2006)) with mean 1, standard deviation < 1 , and autocorrelation > 0 . For applications in this paper, we set the standard deviation to around 0.9 and the autocorrelation coefficient to 0.9, though specific values are noted in the Appendix. The intuition behind this algorithm is similar to parallel tempering (Swendsen and Wang (1986); Geyer (1991); Desjardins et al. (2010b); Brakel et al. (2012); Desjardins et al. (2010a, 2014)). When β is small, the fantasy particles will be able to explore the space quickly. Setting the mean to $\beta = 1$ ensures that the sampled distribution stays *close* to the true distribution, while setting the autocorrelation close to 1 ensures that the inverse temperatures evolve slowly relative to the fantasy particles, which can remain in quasi-equilibrium. Unlike parallel tempering, this *driven sampling* algorithm does not sample from the exact distribution of the RBM. Instead, the driven sampling algorithm samples from a *similar* distribution that has fatter tails (see Figure 3). However, it adds little computational overhead and generally improves training outcomes.

2.7 Using KL Divergences to Monitor Training

We monitor both the forward and reverse KL divergences during training. Following Wang et al. (2009), let $\{X_i\}_{i=1}^n$ and $\{Y_i\}_{i=1}^m$ be samples drawn from densities p and q . Let $\rho_n(i)$ be the distance from X_i to its nearest neighbor in $\{X_j\}_{j \neq i}$, and $\nu_m(i)$ be the distance from

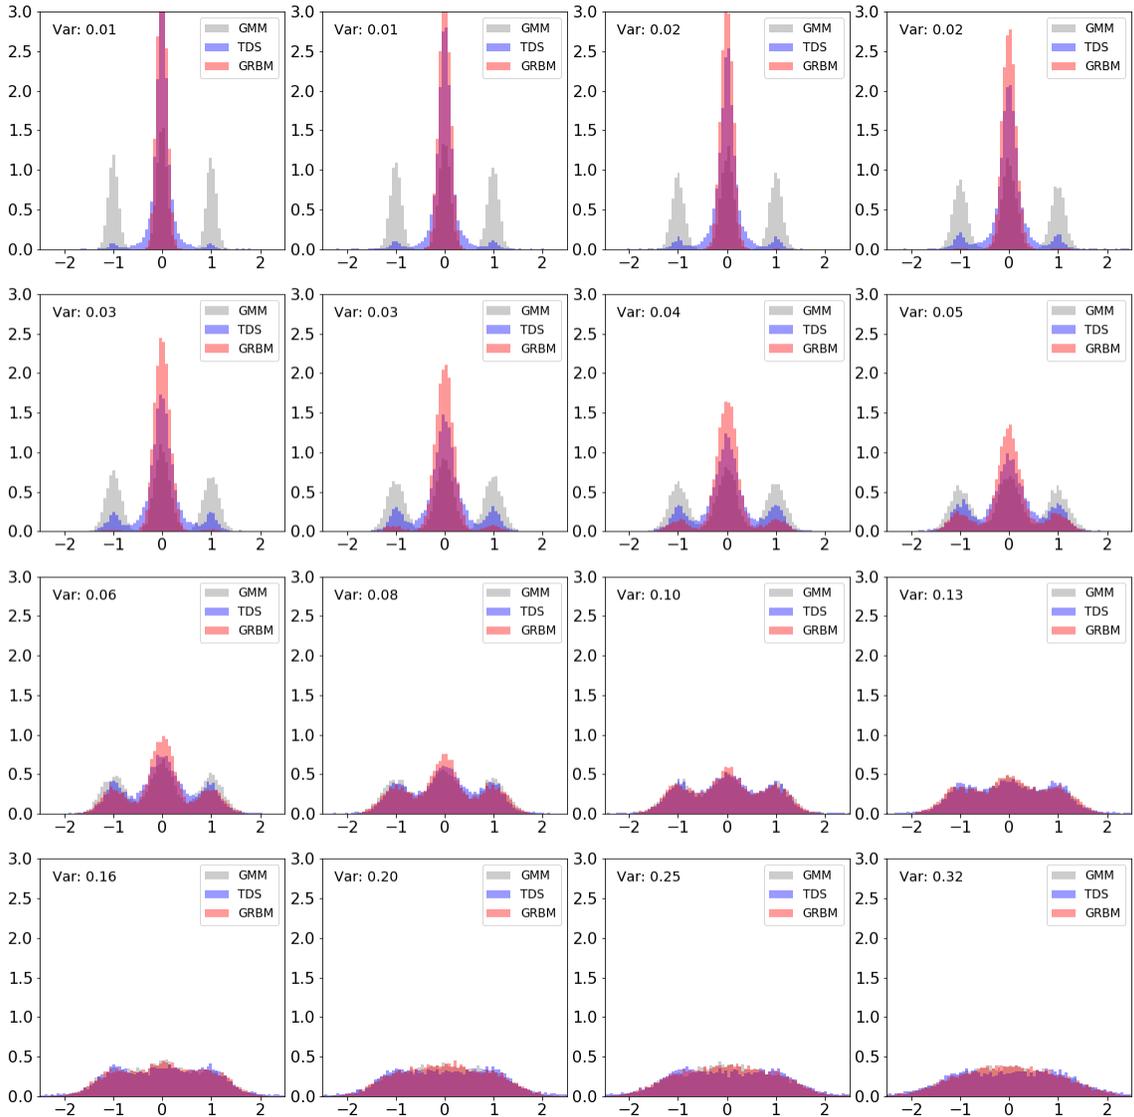


Figure 3: **Sampling with a driven sampler.** Comparison of temperature driven sampling (TDS) to regular Gibbs sampling. The RBMs have a single Gaussian visible layer and a softmax hidden layer with 3 hidden units that encode the modes of a mixture of 3 Gaussians. The standard deviation of the inverse temperature was set to 0.9 for the driven sampler.

X_i to its nearest neighbor in $\{Y_i\}$. Then,

$$D_{KL}(p||q) \approx \frac{d}{n} \sum_{i=1}^n \frac{\nu_m(i)}{\rho_n(i)} + \log \frac{m}{n-1} \quad (17)$$

where d is the dimension of the space (i.e., the number of visible units). The reverse KL divergence can be computed by reversing the identities of X and Y . In practice, we

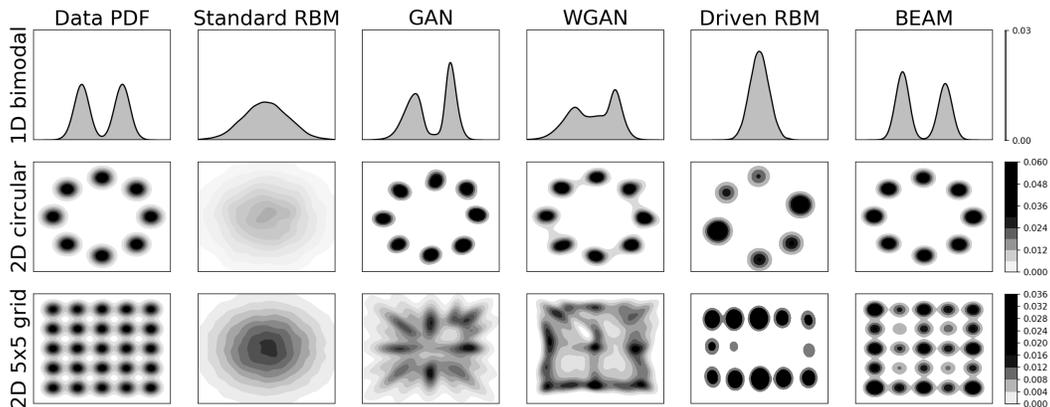


Figure 4: **Comparison of generative models on mixtures of Gaussians.** Three datasets constructed from mixtures of Gaussians: a 1-D mixture of two Gaussians, a 2-D mixture of eight Gaussians arranged in a circle, and a 2-D mixture of Gaussians arranged on a 5x5 grid. Distributions of fantasy particles from a standard RBM, a vanilla GAN, and a Wasserstein GAN (WGAN) are compared to distributions of fantasy particles from a RBM trained with a driven sampler and to a BEAM.

monitor the KL divergences using a held-out validation set consisting of 10% of the data. For computational reasons, we compute the KL divergences on minibatches of the validation set and then average the values.

3. Results

We present empirical results on BEAMs using some datasets that are commonly used to test generative models. We aim to demonstrate four key results:

1. RBMs produce poor results because the reverse KL divergence increases during training even though the forward KL divergence decreases.
2. BEAMs trained with a driven sampler minimize both the forward and reverse KL divergences, leading to better results than RBMs trained by standard methods.
3. BEAMs produce results that are comparable to, or better than, GANs on multiple benchmarks.
4. The simplicity of the adversary ensures that BEAM training is stable.

3.1 Mixture Models

Our first set of experiments are on a series of 1 and 2-dimensional Mixtures of Gaussians (MoGs) similar to those used in the Wasserstein GAN paper (Arjovsky et al. (2017b)). We compare the results from five different generative models. Models from the literature include a vanilla GAN (Goodfellow et al. (2014); greydanus (2017)), a Wasserstein GAN (Arjovsky et al. (2017a,b); Arjovsky (2017)), and a Gaussian-Bernoulli RBM (Cho et al. (2013)). Our

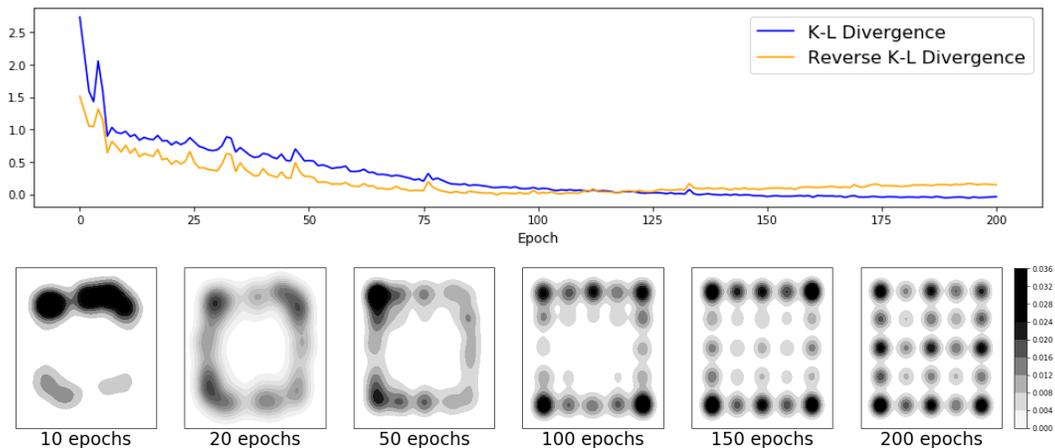


Figure 5: **Training a BEAM on a 2-D Mixture of Gaussians (MoGs) arranged in a 5x5 grid.** Top panel shows estimates of the forward KL divergence, $D_{\text{KL}}(p_d \parallel p_\theta)$, and the reverse KL divergence, $D_{\text{KL}}(p_\theta \parallel p_d)$, per training epoch. Right panels show distributions of fantasy particles at various epochs during training.

models include a Gaussian-Bernoulli RBM trained with the driven sampler and a Gaussian-Bernoulli BEAM with equally weighted likelihood and adversarial losses. All of the RBM based models have the exact same architecture. Details on the model architectures and training parameters are given in the Appendix.

Figure 4 shows a comparison of fantasy particles from each of the generative models along with the corresponding data distributions. A standard RBM trained using persistent contrastive divergence with 100 update steps per gradient evaluation fails to learn that the data distribution has multiple modes. Instead, it spreads the model density across the support of the data distribution. The vanilla GAN and the WGAN are both able to learn the 1-D mixture of two Gaussians and the 2-D mixture of eight Gaussians, but fail on the 2-D MoGs arranged in a 5x5 grid. Surprisingly, our results with the vanilla GAN are significantly better than those reported in the literature Arjovsky et al. (2017b) and are comparable in quality to the results with WGAN. Training the Gaussian-Bernoulli RBM using the driven sampler leads to improvements over the standard RBM. Notably, the BEAM is the only model that learns all three datasets.

Training an RBM as a BEAM decreases both the forward and reverse KL divergences, as shown in Figure 5 for the MoGs arranged in a 5x5 grid. In the early stages of training, the BEAM fantasy particles are spread out across the support of the data distribution – capturing the modes near the edge of the grid. These early epochs resemble the distributions obtained with GANs, which also concentrate density in the modes near the edge of the grid. As training progresses, the BEAM progressively learns to capture the modes near the center of the grid.

3.2 MNIST

The MNIST dataset of handwritten images (LeCun and Cortes (2010)) is one of the most widely used benchmarks in machine learning. We present results on MNIST with binary (black and white) images and with continuous (grayscale) images.

3.2.1 BINARY MNIST

Most work on RBMs has focused on learning distributions of binary data using Bernoulli visible layers. Therefore, we present our results on binary MNIST before moving on to continuous MNIST. We compare three different generative models on binary MNIST: a Bernoulli-Bernoulli RBM, a Bernoulli-Bernoulli RBM trained using temperature driven sampling, and a Bernoulli-Bernoulli BEAM. Details of the architectures and training hyperparameters are provided in the Appendix.

One thing of note is that we train the BEAM in two phases. The critic in a BEAM uses the hidden unit activities as features, but these features are not useful during the early stages of training when there is little mutual information between the visible and hidden units of the generator. Therefore, we use regular maximum likelihood training with persistent contrastive divergence and driven sampling for the first 25 epochs. After 25 epochs, we change the relative weights of the likelihood and the adversary in the loss function to $\gamma = 0.1$ so that the adversarial term dominates the gradient and train for an additional 30 epochs.

Our Bernoulli-Bernoulli RBM generates reasonable looking fantasy particles for the binary MNIST dataset (Figure 6), in agreement with previous work by Tieleman (2008). Still, the Bernoulli-Bernoulli BEAM generates better and more diverse fantasy particles. This is because the RBM and BEAM achieve roughly the same forward KL divergence, but the BEAM has a much lower reverse KL divergence (Table 1). Moreover, comparing the results to the driven RBM shows that this is clearly not only due to the use of the temperature driven sampler.

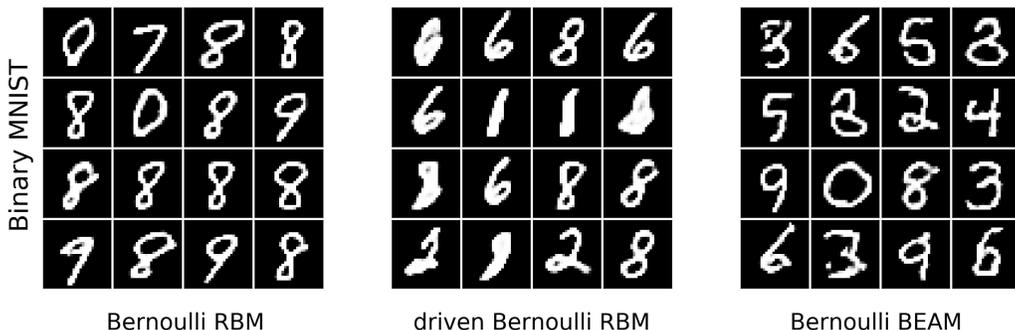


Figure 6: **Comparison of binary MNIST fantasy particles.** Sixteen particles sampled at random from each of the generators. The RBM fantasy particles were randomly initialized and sampled for 100 MCMC steps; the figure shows $\langle \mathbf{v} \rangle_{p_{\theta}(\mathbf{v}|\mathbf{h})}$ computed from the last iteration.

| Model | Forward KL | Reverse KL |
|----------------------|------------|------------|
| Bernoulli RBM | 67.47 | 168.37 |
| Driven Bernoulli RBM | 96.84 | 192.68 |
| Bernoulli BEAM | 70.48 | 63.58 |

Table 1: Performance on binary MNIST at the end of training.

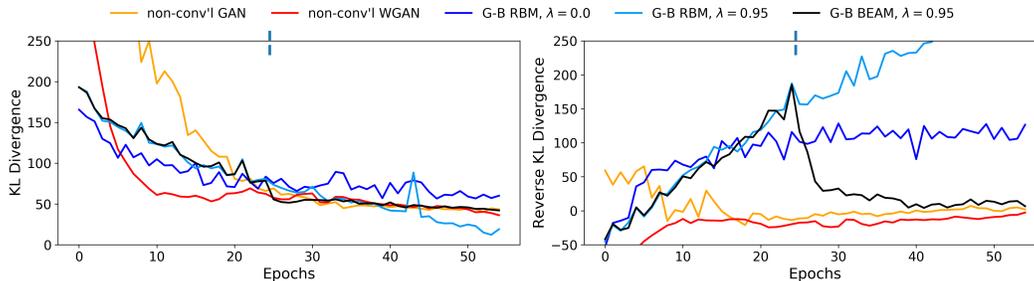


Figure 7: **Training metrics on continuous MNIST.** The forward KL divergence, $D_{\text{KL}}(p_d \| p_\theta)$, and the reverse KL divergence, $D_{\text{KL}}(p_\theta \| p_d)$ divergence during training on MNIST. Adversarial training for the BEAM begins after epoch 25 (vertical dashed line).

We do not show any GANs for the binary MNIST problem. In general, it is difficult to train GANs on discrete data due to the inability to backprop through a discrete variable (though, there are ways around this problem as in Yu et al. (2017)). Thus, one advantage of a BEAM is that it is much easier to train on discrete data than a GAN and much easier to train on continuous data than a standard RBM.

3.2.2 CONTINUOUS MNIST

Compared to binary MNIST, continuous MNIST is a much more difficult problem for an RBM (Yamashita et al. (2014)) but it is an easier problem for a GAN. Here, we compare three different RBM based models: a Gaussian-Bernoulli RBM, a Gaussian-Bernoulli RBM

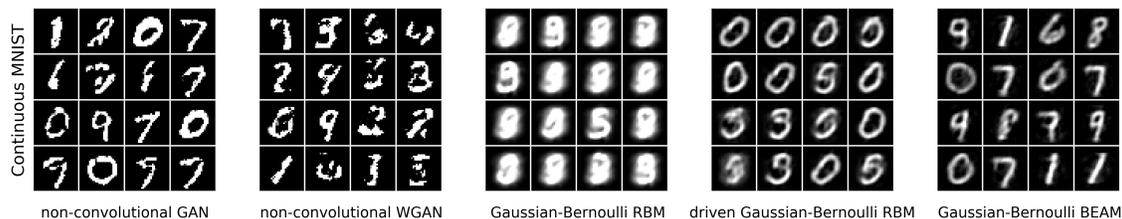


Figure 8: **Comparison of continuous MNIST fantasy particles.** Sixteen particles sampled at random from each of the generators. The RBM fantasy particles were randomly initialized and sampled for 100 MCMC steps; the figure shows $\langle \mathbf{v} \rangle_{p_\theta(\mathbf{v}|\mathbf{h})}$ computed from the last iteration.

| Model | Forward KL | Reverse KL |
|---------------------|------------|------------|
| GAN | 43.40 | 1.93 |
| WGAN | 36.17 | -2.65 |
| Gaussian RBM | 60.09 | 126.52 |
| Driven Gaussian RBM | 19.05 | 347.67 |
| Gaussian BEAM | 42.10 | 6.70 |

Table 2: Performance on continuous MNIST at the end of training.

with a temperature driven sampler, and a Gaussian-Bernoulli BEAM. We also compare to a fully connected GAN and a fully connected WGAN.

It is important to note that none of these models is designed to produce state-of-the-art results on MNIST; for example, you can get better results using convolutional, rather than fully-connected, networks (Figure 13). However, restricting the analyses to the chosen architectures provides a cleaner comparison of the different training approaches.

RBM based architectures trained by maximum likelihood will decrease the forward KL divergence. This is shown clearly in Figure 7 – the forward KL divergence decreases during training of the Gaussian-Bernoulli RBM, the Gaussian-Bernoulli RBM with a driven sampler, and the Gaussian-Bernoulli BEAM. However, the figure also clearly shows that the reverse KL divergence increases during training. The training metrics for the BEAM rapidly diverge from the RBMs once the adversary is turned on after epoch 25. The reverse KL divergence of the BEAM quickly drops towards zero while the reverse KL divergence of the RBMs continue to rise. By the end of training, the BEAM obtains comparable, or better, metrics than all other architectures (Table 2).

Fantasy particles for continuous MNIST are shown in the top row of Figure 8. The non-convolutional GAN, non-convolutional WGAN, and the BEAM have similar metrics at the end of training. The errors that they make, however, are qualitatively different. The GANs produce sharp images that are a bit blotchy, whereas the BEAM produces smooth images that are a bit blurry. The regular Gaussian-Bernoulli RBM fails to produce reasonable digits at all, whereas the Gaussian-Bernoulli RBM trained with the driven sampler is a bit better.

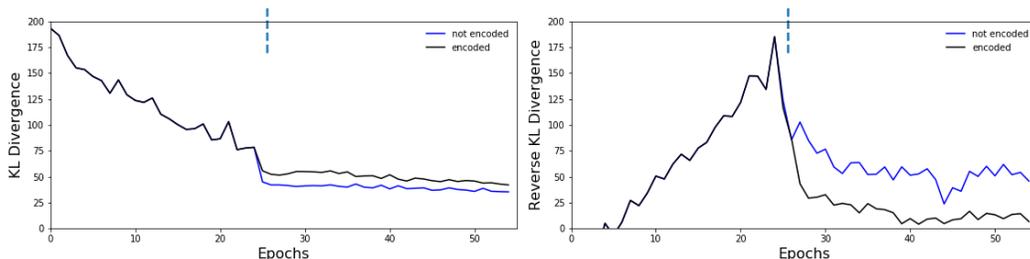


Figure 9: **Comparing a BEAM with the critic on the hidden layer to one with the critic on the visible layer.** The KL divergences of two BEAMs trained on MNIST differing only in whether or not the critic acts on the encoded data or directly on the visible data. Adversarial training begins after 25 epochs.

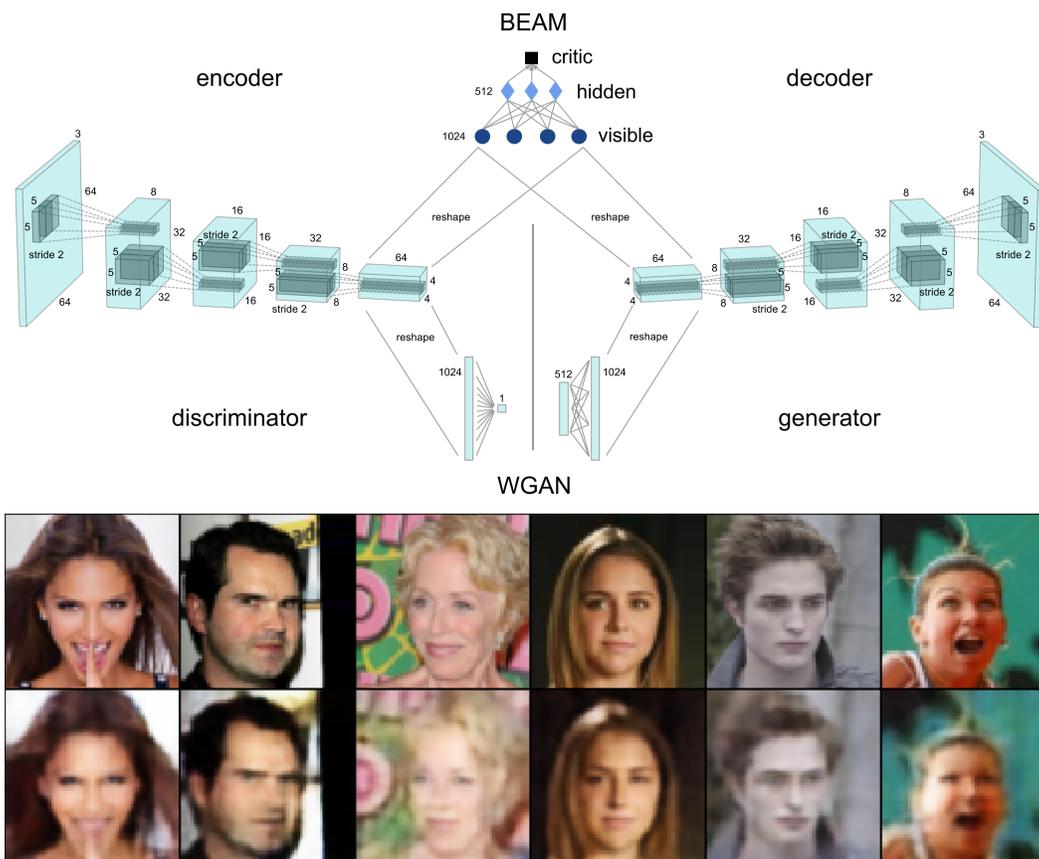


Figure 10: **Architectures used in the CelebA experiment.** The autoencoder is purely convolutional, with the encoder (decoder) shown by the stacks of convolutional (deconvolutional) layers. The BEAM uses the flattened encoded features as the visible units. The GAN/WGAN uses the same encoder (decoder) architectures for the discriminator (generator), with added single fully connected layers. The autoencoder, BEAM, and GANs are each trained fully independently. Table 5 lays out the parameters of the models. Example reconstructions from the autoencoder are shown in the last row.

Throughout, we have presented BEAMs as an adversarial approach to training RBMs where the hidden unit activities of the RBM are used as features for the critic. We claim that this allows us to use a simple classifier to approximate the optimal critic. However, it is possible to train an RBM against an adversary that uses the visible units directly (as in a GAN). Empirically, we have found that applying the critic to the hidden unit activities works better; see Figure 9.

3.3 Celebrity Faces

Natural images present a more complex dataset for which model performance can be easily determined. We use the CelebA dataset, consisting of $64 \times 64 \times 3$ pictures of celebrities' faces, to demonstrate that BEAMs scale to more complex problems. This dataset requires

convolutional architectures to obtain good performance. Because exploring convolutional RBMs is orthogonal to the purpose of this work, we use a separately trained convolutional autoencoder to extract features from the images (Figure 10). These features are used as the visible input to the BEAM. The autoencoder is trained with sufficient depth and number of features to provide high-quality reconstructions of the data (Figure 10).

As with the MNIST examples, we train the CelebA BEAM in two phases. For the first 15 epochs, we use the log-likelihood objective function with persistent contrastive divergence and driven sampling. After this phase we train for an additional 50 epochs using the combined log-likelihood-adversarial objective function in Equation 9 with $\gamma = 0.1$.

For comparison, we train a DCGAN and DCWGAN using the same convolutional architecture as the autoencoder that was used to extract features for BEAM training. That is, the DCGAN/DCWGAN generator uses an initial fully connected layer followed by the same architecture as the decoder of the autoencoder and the DCGAN/DCWGAN critic uses the same architecture as the encoder of the autoencoder followed by a fully connected layer to a single unit. The DCGAN/DCWGAN share many of architectural features with the autoencoder, but do not share any parameters. Instead, the DCGAN and the DCWGAN were trained end-to-end on CelebA.

Images are generated from the BEAM by sampling fantasy particles and passing them through the decoder. Example generated images from the BEAM, DCGAN, and DCWGAN are shown in Figure 11. It is clear that the BEAM images are internally consistent with clear features across each face. However, the images are a bit blurry – especially towards the corners of the image in the backgrounds. The images produced by the DCGAN and DCWGAN have sharper local features, but notably poorer global correlations. Although the images produced by the GANs are not particularly high-quality, they are qualitatively similar to results appearing in the literature. To illustrate this, we have directly reproduced fantasy images from a DCWGAN that were published in Li et al. (2017) (see Figure 11d).

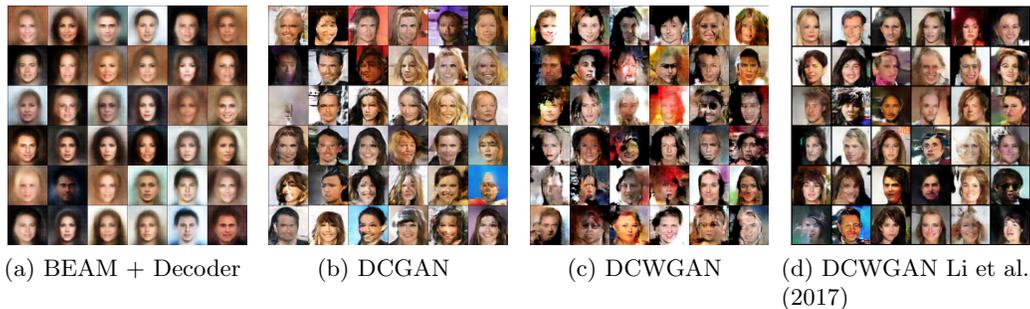


Figure 11: **Comparison of CelebA fantasy particles.** 36 fantasy images sampled from the (a) BEAM with the decoder network, (b) a deep convolutional GAN, and (c) a deep convolutional WGAN. Our implementations were chosen so that each model has very similar architectures. These architectures and training parameters are provided in the Appendix. For comparison, we directly reproduce the CelebA results of a deep convolutional WGAN as reported by Li et al. (2017).

We note that it is possible to obtain sharper features from the BEAM at the expense of less consistent images by optimizing training to lower the forward KL divergence at the expense of an increased reverse KL divergence. Furthermore, additional approaches such as centering layers and using deep models produce notably better images; see Figure 14 for an example using centered layers.

4. Discussion

We have introduced a novel formulation of RBMs, called BEAMs, that uses an adversary acting on the hidden unit activations from the RBM to supplement the traditional likelihood-based training. The additional adversarial loss term ensures that training minimizes both the forward and reverse KL divergences, allowing the model to accurately learn distributions with multiple modes. We have shown that BEAMs excel at a variety of applications, outperforming GANs that use significantly larger computational budgets.

As the machine learning community increasingly turns its attention to unsupervised learning problems, it is valuable to place this work into a larger context. The deep learning revolution has driven tremendous advances on supervised learning problems, and a primary outcome is that feed-forward neural networks have become a powerful tool. GANs and variational autoencoders can be thought of as a natural extension of the broad learning capacity of neural networks and the flexibility of backpropagation, and are tools of choice in many applications. This is further supported by the software ecosystem for machine learning, which makes many sophisticated tools easily accessible.

RBMs played an active role in kicking off the deep learning revolution (Hinton and Salakhutdinov (2006)), but their development slowed with the increased focus on supervised learning and a general attitude that they were unsuited to more complex problems. However, there are reasons to believe that RBMs will be fundamental in advancing unsupervised learning:

- Novel training algorithms and novel model architectures can dramatically improve performance.
- RBMs have several analytic handles to understand models and develop training strategies.
- Increased capacity to handle complex datasets can be developed through a progressive set of challenging applications.

We hope this work reinforces the promise of RBMs.

Acknowledgments

We would like to thank Pankaj Mehta and David Schwab for comments on an early draft of this manuscript.

Appendix A.

In a general adversarial approach to learning, we train a Boltzmann machine, $p_\theta(\mathbf{v}, \mathbf{h})$, to minimize a compound objective function $\mathcal{C} = -\gamma\mathcal{L} + (1 - \gamma)\mathcal{A}$. The compound objective function represents a tradeoff between maximum likelihood learning ($\gamma = 1$) and adversarial learning ($\gamma = 0$). Just as with maximum likelihood, the compound objective function can be optimized using stochastic gradient descent. Using a compound objective function helps to mitigate some of the instability problems that plague traditional GANs. For example, the gradient does not vanish even if the discriminator is completely untrained because there is always the term from the likelihood.

We need to compute the derivatives of the compound objective function in order to minimize it. The differential operator is linear, so we can distribute it across the two terms:

$$\partial_\theta \mathcal{C} = \gamma \partial_\theta \mathcal{L} - (1 - \gamma) \partial_\theta \mathcal{A}. \quad (18)$$

The first term can be computed from Equation 2 (Main Text). So, all we need to do is compute the second term. It turns out that derivatives of this form can be computed using a simple formula when the model is a Boltzmann machine.

Let $T(\mathbf{v}, \mathbf{h})$ be a critic function and

$$\mathcal{A} := \int d\mathbf{v} d\mathbf{h} p(\mathbf{v}, \mathbf{h}) T(\mathbf{v}, \mathbf{h}) \quad (19)$$

be the associated adversarial loss. This formulation reduces to the adversarial loss for a BEAM when $T(\mathbf{v}, \mathbf{h}) = T(\mathbf{h})$ is independent of the visible units, but we derive the general case. We need to compute $\partial_\theta \mathbb{E}_{p_\theta(\mathbf{v}, \mathbf{h})}[T(\mathbf{v}, \mathbf{h})]$. First, we use the stochastic derivative trick:

$$\begin{aligned} \partial_\theta \mathcal{A} &= \frac{\partial}{\partial \theta} \int d\mathbf{v} d\mathbf{h} p(\mathbf{v}, \mathbf{h}) T(\mathbf{v}, \mathbf{h}) \\ &= \int d\mathbf{v} d\mathbf{h} T(\mathbf{v}, \mathbf{h}) \frac{\partial}{\partial \theta} p(\mathbf{v}, \mathbf{h}) \\ &= \int d\mathbf{v} d\mathbf{h} T(\mathbf{v}, \mathbf{h}) \frac{p(\mathbf{v}, \mathbf{h})}{p(\mathbf{v}, \mathbf{h})} \frac{\partial}{\partial \theta} p(\mathbf{v}, \mathbf{h}) \\ &= \int d\mathbf{v} d\mathbf{h} T(\mathbf{v}, \mathbf{h}) p(\mathbf{v}, \mathbf{h}) \partial_\theta \log p(\mathbf{v}, \mathbf{h}) \\ &= \langle T(\mathbf{v}, \mathbf{h}) \partial_\theta \log p(\mathbf{v}, \mathbf{h}) \rangle_{p(\mathbf{v}, \mathbf{h})} \end{aligned} \quad (20)$$

We can write the model distribution of a Boltzmann machine as $p_\theta(\mathbf{v}, \mathbf{h}) = Z_\theta^{-1} e^{-E_\theta(\mathbf{v}, \mathbf{h})}$ so that $\log p_\theta(\mathbf{v}, \mathbf{h}) = -E_\theta(\mathbf{v}, \mathbf{h}) - \log Z_\theta$, with $Z_\theta = \int d\mathbf{v} d\mathbf{h} e^{-E_\theta(\mathbf{v}, \mathbf{h})}$. Then, we have $\partial_\theta \log p_\theta(\mathbf{v}, \mathbf{h}) = -\langle -\partial_\theta E_\theta(\mathbf{v}, \mathbf{h}) \rangle_{p_\theta(\mathbf{v}, \mathbf{h})} - \partial_\theta E_\theta(\mathbf{v}, \mathbf{h})$. Plugging this in, we find:

$$\partial_\theta \mathcal{A} = -\langle T(\mathbf{v}, \mathbf{h}) \rangle_{p_\theta(\mathbf{v}, \mathbf{h})} \langle -\partial_\theta E_\theta(\mathbf{v}, \mathbf{h}) \rangle_{p_\theta(\mathbf{v}, \mathbf{h})} + \langle T(\mathbf{v}, \mathbf{h}) (-\partial_\theta E_\theta(\mathbf{v}, \mathbf{h})) \rangle_{p_\theta(\mathbf{v}, \mathbf{h})} \quad (21)$$

$$= \text{Cov}_{p_\theta(\mathbf{v}, \mathbf{h})}[T(\mathbf{v}, \mathbf{h}), -\partial_\theta E_\theta(\mathbf{v}, \mathbf{h})]. \quad (22)$$

Appendix B.

Our approach to accelerated sampling, which we call Temperature Driven Sampling (TDS), greatly improves the ability to train Boltzmann machines without incurring significant additional computational cost. The algorithm is a variant of a sequential Monte Carlo sampler.

A collection of m samples are evolved independently using Gibbs sampling updates from the model. Note that this is not the same as running multiple chains for a parallel tempering algorithm because each of the m samples in the sequential Monte Carlo sampler will be used compute statistics, as opposed to just the samples from the $\beta = 1$ chain during parallel tempering. Each of these samples has an inverse temperature that is drawn from a distribution with mean $E[\beta] = 1$ and variance $\text{Var}[\beta] < 1$. The inverse temperatures of each sample are independently updated once for every Gibbs sampling iteration of the model; however, the updates are autocorrelated across time so that the inverse temperatures are slowly varying. As a result, the collection of samples are drawn from a distribution that is *close to* the model distribution, but with fatter tails. This allows for much faster mixing, while ensuring that the model averages (computed over the collection of m samples) remain close approximations to averages computed from the model with $\beta = 1$.

Input:

Autocorrelation coefficient $0 \leq \phi < 1$.

Variance of the distribution $\text{Var}[\beta] < 1$.

Current value of β .

Set: $\nu = 1/\text{Var}[\beta]$ and $c = (1 - \phi)\text{Var}[\beta]$.

Draw $z \sim \text{Poisson}(\beta * \phi/c)$.

Draw $\beta' \sim \text{Gamma}(\nu + z, c)$.

return β'

Algorithm 1: Sampling from an autocorrelated Gamma distribution.

Details of the TDS algorithm are provided in Algorithms 1 and 2. Note that this algorithm includes a standard Gibbs sampling based sequential Monte Carlo sampler in the limit that $\text{Var}[\beta] \rightarrow 0$. The samples drawn with the TDS algorithm are *not* samples from the equilibrium distribution of the Boltzmann machine. In principle, it is possible to reweight these samples to correct for the bias due to the varying temperature. In practice, we have not found that reweighting is necessary.

Input:

Number of samples m .

Number of update steps k .

Autocorrelation coefficient for the inverse temperature $0 \leq \phi < 1$.

Variance of the inverse temperature $\text{Var}[\beta] < 1$.

Initialize:

Randomly initialize m samples $\{(\mathbf{v}_i, \mathbf{h}_i)\}_{i=1}^m$.

Randomly initialize m inverse temperatures $\beta_i \sim \text{Gamma}(1/\text{Var}[\beta], \text{Var}[\beta])$.

for $t = 1, \dots, k$ **do**

| | |
|--|---|
| for $i = 1, \dots, m$ do | Update β_i using Algorithm 1. |
| | Update $(\mathbf{v}_i, \mathbf{h}_i)$ using Gibbs sampling. |
| end | |

end

Algorithm 2: Temperature Driven Sampling.

Appendix C.

This section provides some additional experimental details and results.

Gaussian Mixtures

Table 3 lays out the parameters of the Gaussian mixture comparison examples. It is interesting to note just how few parameters are required by the BEAM to model this data.

MNIST

Table 4 lays out the parameters of the models for the MNIST examples.

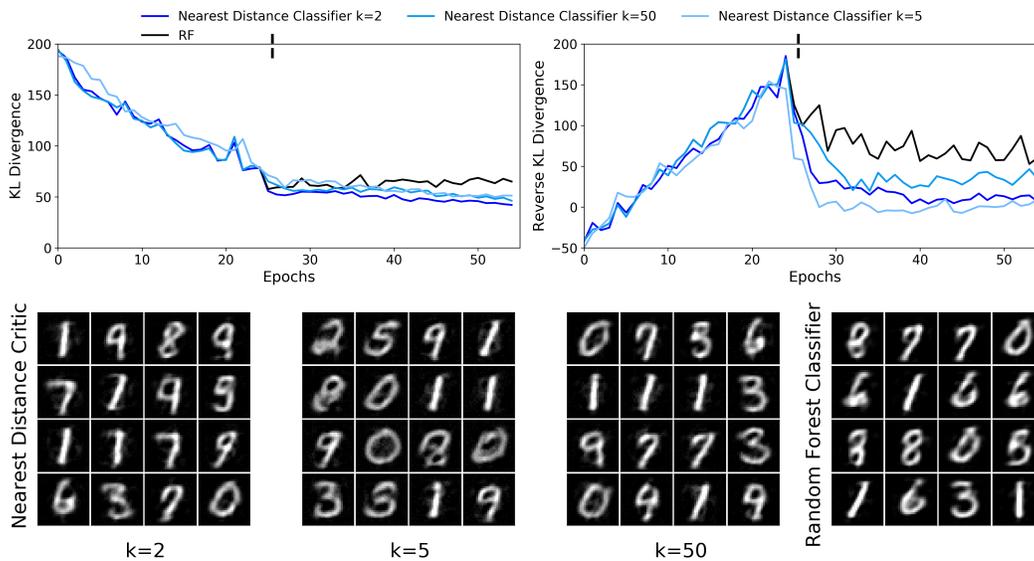


Figure 12: Gaussian-Bernoulli BEAMs trained on continuous MNIST with different critics.

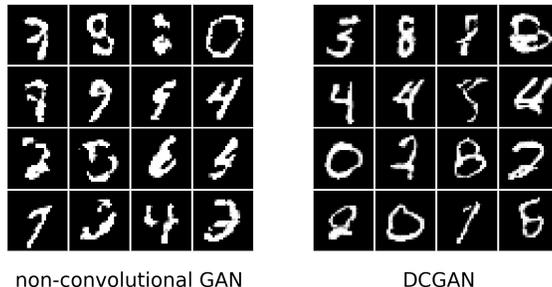


Figure 13: Comparison of a non-convolutional GAN to a DCGAN on continuous MNIST.

Celebrity Faces

There is plenty of room to improve the quality of generated faces by employing more advanced RBM training techniques. For example, centering the RBMs tends to improve the variance in the generated faces and increases the definition of the hair and face edges.

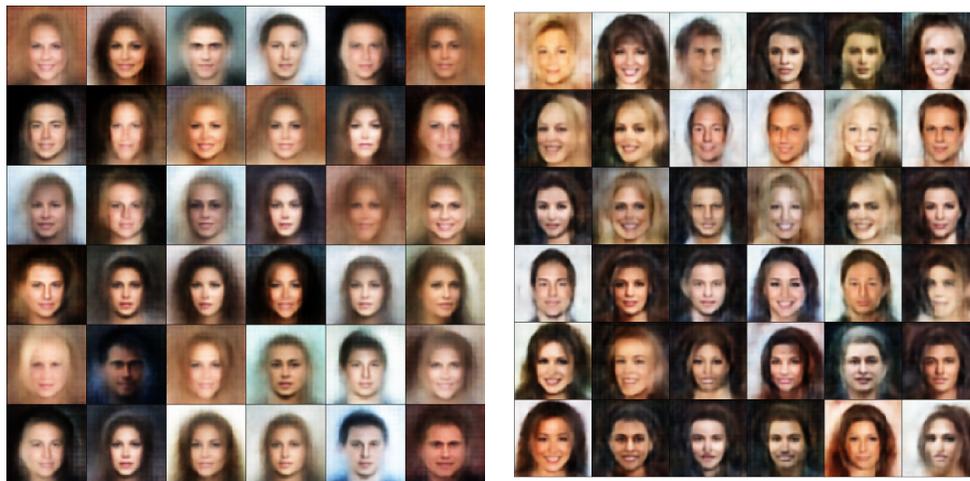


Figure 14: **BEAM vs. centered BEAM fantasy particles** Example fantasy particles generated by a BEAM using centered visible layer.

| | | | | |
|-------------------------|--|-------------------------|--------|--------|
| Bimodal Gaussian | 10 ⁴ samples, batch size 100 | | | |
| GAN/WGAN | fully-connected with ReLU activations, WGAN weight clamping .1 | | | |
| | generator dimensions | critic dimensions | epochs | lr |
| | 1 – 32 | 1 – 32 – 1 | 100 | 0.002 |
| RBM/BEAM | distance-weighted nearest-neighbor critic $k = 5$, $\lambda = 0.5$ for BEAM | | | |
| | dims | MCMC steps | epochs | lr |
| | 1 – 10 | 100 | 10 | .1 |
| Radial Gaussian | 10 ⁴ samples, batch size 100 | | | |
| GAN/WGAN | fully-connected with ReLU activations, WGAN weight clamping .1 | | | |
| | generator dimensions | critic dimensions | epochs | lr |
| | 2 – 64 – 64 | 2 – 64 – 64 – 1 | 300 | 0.001 |
| RBM/BEAM | distance-weighted nearest-neighbor critic $k = 5$, $\lambda = 0.5$ for BEAM | | | |
| | dims | MCMC steps | epochs | lr |
| | 1 – 10 | 100 | 30 | .15 |
| Grid Gaussian | 10 ⁵ samples, batch size 1000 | | | |
| GAN/WGAN | fully-connected with ReLU activations, WGAN weight clamping .1 | | | |
| | generator dimensions | critic dimensions | epochs | lr |
| | 2 – 128 – 128 – 128 | 2 – 128 – 128 – 128 – 1 | 300 | .00001 |
| RBM/BEAM | distance-weighted nearest-neighbor critic $k = 5$, $\lambda = 0.5$ for BEAM | | | |
| | dims | MCMC steps | epochs | lr |
| | 1 – 20 | 100 | 200 | .2 |

Table 3: **Gaussian mixture architectures and hyperparameters** All GAN/WGAN models use ReLU activations between fully-connected layers. Network weights are initialized with normal distributions of standard deviation 0.2, with biases zero-initialized. The beta standard deviation for the driven sampler is set to 0 for RBM, .9 for driven RBM and BEAMs. The RBMs’ learning rates decrease according to a power-law decay, and all training uses ADAM optimization with beta = (0.5, 0.9).

| MNIST | $6 \cdot 10^5$ samples, batch size 100 | | | |
|----------|---|---------------------|----------------|---------------------|
| GAN/WGAN | fully-connected with ReLU activations, sigmoid on generator, WGAN weight clamping .01 | | | |
| | generator dimensions | critic dimensions | epochs | lr |
| | 100 – 164 – 164 | 784 – 164 – 164 – 1 | 55 | .001 |
| RBM/BEAM | distance-weighted nearest-neighbor critic $k = 2$, $\lambda = 0.1$ for BEAM | | | |
| | dims | MCMC steps | epochs | lr |
| | 784 – 200 | 5 | 25 ML, 30 adv. | .001 ML, .0001 adv. |

Table 4: **Gaussian mixture architectures and hyperparameters** All GAN/WGAN models use ReLU activations between fully-connected layers. Generator and discriminator weights are initialized with normal distributions of standard deviation 0.1 and 0.02 resp., with biases zero-initialized. All training uses ADAM optimization with beta = (0.5, 0.9) for the GANs and (0.9, 0.999) for the BEAM. For the BEAM, the beta standard deviation for the driven sampler is set to .95. The RBMs’ learning rates decrease according to a power-law decay.

| CelebA | 202599 samples, batch size 128 | | | |
|---------------|--|---|----------------|------------------------------|
| GAN/WGAN | conv. with batch-norm and ReLU activations, WGAN weight clamping .01 | | | |
| | generator dimensions | critic dimensions | epochs | lr |
| | $64 \times 4^2 - 64 \times 8^2$ | $64 \times 16^2 - 32 \times 32^2 - 3 \times 64^2$ | 10 | .0002 |
| BEAM | distance-weighted nearest-neighbor critic $k = 5$, $\lambda = 0.1$ for BEAM | | | |
| | dims | MCMC steps | epochs | lr |
| | 1024 – 512 | 100 | 10 ML, 40 adv. | 10^{-3} ML, 10^{-4} adv. |

Table 5: **CelebA architectures and hyperparameters** All GAN/WGAN models use ReLU activations between fully-connected layers. Generator and discriminator weights are initialized with normal distributions of standard deviation 0.02, with biases zero-initialized. The beta standard deviation for the driven sampler is set to 0 for RBM, .95 for the BEAM. All training uses ADAM optimization with beta = (0.5, 0.9) for the GANs and (0.9, 0.999) for the BEAM. The RBMs’ learning rates decrease according to a power-law decay.

References

- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- Tarik Arici and Asli Celikyilmaz. Associative adversarial networks. *arXiv preprint arXiv:1611.06953*, 2016.
- Martin Arjovsky. Wasserstein gan. <http://github.com/martinarjovsky/WassersteinGAN>, 2017.
- Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017a.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 214–223, 2017b.
- Philemon Brakel, Sander Dieleman, and Benjamin Schrauwen. Training restricted boltzmann machines with multi-tempering: harnessing parallelization. *Artificial Neural Networks and Machine Learning–ICANN 2012*, pages 92–99, 2012.
- Kyung Hyun Cho, Tapani Raiko, and Alexander Ilin. Gaussian-bernoulli deep boltzmann machine. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–7. IEEE, 2013.
- Guillaume Desjardins, Aaron Courville, Yoshua Bengio, Pascal Vincent, and Olivier Delalleau. Parallel tempering for training of restricted boltzmann machines. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 145–152. MIT Press Cambridge, MA, 2010a.
- Guillaume Desjardins, Aaron Courville, Yoshua Bengio, Pascal Vincent, and Olivier Delalleau. Tempered markov chain monte carlo for training of restricted boltzmann machines. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 145–152, 2010b.
- Guillaume Desjardins, Heng Luo, Aaron Courville, and Yoshua Bengio. Deep tempering. *arXiv preprint arXiv:1410.0123*, 2014.
- Charles J Geyer. Markov chain monte carlo maximum likelihood. 1991.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- Christian Gouriéroux and Joann Jasiak. Autoregressive gamma processes. *Journal of Forecasting*, 25(2):129–152, 2006.

- greydanus. Mnist gan. <http://github.com/greydanus/mnist-gan>, 2017.
- Geoffrey Hinton. A practical guide to training restricted boltzmann machines. *Momentum*, 9(1):926, 2010.
- Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Training*, 14(8), 2006.
- Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.
- Geoffrey E Hinton and Ruslan R Salakhutdinov. A better way to pretrain deep boltzmann machines. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2012.
- Geoffrey E Hinton and Terrence Joseph Sejnowski. *Unsupervised learning: foundations of neural computation*. MIT press, 1999.
- Artur Kadurin, Sergey Nikolenko, Kuzma Khrabrov, Alex Aliper, and Alex Zhavoronkov. drugan: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico. *Molecular Pharmaceutics*, 14(9):3098–3104, 2017.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Volodymyr Kuleshov and Stefano Ermon. Neural variational inference and learning in undirected graphical models. In *Advances in Neural Information Processing Systems*, pages 6737–6746, 2017.
- Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- Nicolas Le Roux and Yoshua Bengio. Representational power of restricted boltzmann machines and deep belief networks. *Neural computation*, 20(6):1631–1649, 2008.
- Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd gan: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*, pages 2200–2210, 2017.
- Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate bayesian inference. *arXiv preprint arXiv:1704.04289*, 2017.
- Grégoire Montavon, Klaus-Robert Müller, and Marco Cuturi. Wasserstein training of restricted boltzmann machines. In *Advances in Neural Information Processing Systems*, pages 3718–3726, 2016.

- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pages 271–279, 2016.
- Jason Tyler Rolfe. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- Ruslan Salakhutdinov and Geoffrey Hinton. Deep boltzmann machines. In *Artificial Intelligence and Statistics*, pages 448–455, 2009.
- Nitish Srivastava and Ruslan R Salakhutdinov. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pages 2222–2230, 2012.
- Ilya Sutskever, Geoffrey E Hinton, and Graham W Taylor. The recurrent temporal restricted boltzmann machine. In *Advances in Neural Information Processing Systems*, pages 1601–1608, 2009.
- Robert H Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Physical Review Letters*, 57(21):2607, 1986.
- Graham W Taylor and Geoffrey E Hinton. Factored conditional restricted boltzmann machines for modeling motion style. In *Proceedings of the 26th annual international conference on machine learning*, pages 1025–1032. ACM, 2009.
- Graham W Taylor, Geoffrey E Hinton, and Sam T Roweis. Modeling human motion using binary latent variables. In *Advances in neural information processing systems*, pages 1345–1352, 2007.
- Tijmen Tieleman. Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning*, pages 1064–1071. ACM, 2008.
- Qing Wang, Sanjeev R Kulkarni, and Sergio Verdú. Divergence estimation for multidimensional densities via k -nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55(5):2392–2405, 2009.
- Takayoshi Yamashita, Masayuki Tanaka, Eiji Yoshida, Yuji Yamauchi, and Hironobu Fujiyoshi. To be bernoulli or to be gaussian, for a restricted boltzmann machine. In *Pattern Recognition (ICPR), 2014 22nd International Conference on*, pages 1520–1525. IEEE, 2014.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, pages 2852–2858, 2017.