
Learning the Stein Discrepancy for Training and Evaluating Energy-Based Models without Sampling

Will Grathwohl¹ Kuan-Chieh Wang¹ Jörn-Henrik Jacobsen¹ David Duvenaud¹ Richard Zemel¹

Abstract

We present a new method for evaluating and training unnormalized density models. Our approach only requires access to the gradient of the unnormalized model’s log-density. We estimate the Stein discrepancy between the data density $p(x)$ and the model density $q(x)$ defined by a vector function of the data. We parameterize this function with a neural network and fit its parameters to maximize the discrepancy. This yields a novel goodness-of-fit test which outperforms existing methods on high dimensional data. Furthermore, optimizing $q(x)$ to minimize this discrepancy produces a novel method for training unnormalized models which scales more gracefully than existing methods. The ability to both learn and compare models is a unique feature of the proposed method.

1. Introduction

Energy-Based Models (EBMs), also known as unnormalized density models, are perhaps the most flexible way to parameterize a density. They hinge on the observation that any density $p(x)$ can be expressed as

$$p(\mathbf{x}) = \frac{\exp(-E(\mathbf{x}))}{Z}, \quad (1)$$

where $E : \mathbb{R}^D \rightarrow \mathbb{R}$, known as the *energy function*, maps each point to a scalar, and $Z = \int_{\mathbf{x}} \exp(-E(\mathbf{x}))$ is the normalizing constant.

A major benefit of EBMs is that they allow maximal freedom in designing the energy function E . This makes it straightforward to incorporate prior knowledge about the problem, such as symmetries or domain-specific design choices, into the structure of the model. This has

¹University of Toronto and Vector Institute, Toronto, Canada. Correspondence to: Will Grathwohl <wgrathwohl@cs.toronto.edu>.




Figure 1. Density models trained with approximate MCMC samplers can fail to match the data density while still generating high-quality samples. Samples from approximate MCMC samplers follow a *different* distribution than the density they are applied to. It is this induced distribution which is trained to match the data. In contrast, our approach LSD directly matches the model density to the data density without reliance on a sampler.

made EBMs an appealing candidate for applications in physics (Noé et al., 2019), biology (Ingraham et al., 2019), neuroscience (Scellier & Bengio, 2017), and computer vision (LeCun et al., 2007; Osadchy et al., 2007; Xie et al., 2016; 2019; 2018), to name a few.

Despite their many benefits, EBMs present a central challenge which complicates their use: because we cannot efficiently compute the normalizing constant, we cannot compute likelihoods under our model, making training and evaluation difficult. Much prior work on EBMs has relied on MCMC sampling techniques to estimate the likelihood (for evaluation) and its gradient (for training). Other approaches train EBMs by finding easier-to-compute surrogate objectives which have similar optima to the maximum likelihood objective. These include Score Matching (Hyvärinen, 2005) and Noise-Contrastive Estimation (Gutmann & Hyvärinen, 2010).

These original sampling- and score-based approaches were not able to scale to large, high-dimensional datasets as well as subsequently developed alternative models, such as Variational Autoencoders (VAEs) (Kingma & Welling, 2013) and Normalizing Flows (NFs) (Rezende & Mohamed, 2015). These approaches offer more easily scalable training, evaluation, and sampling, but do so at the cost of a more restrictive model parameterization which can lead to well-known prob-

lems like posterior collapse in VAEs (Lucas et al., 2019), and the inability of NFs to model distributions with certain topological structures (Falorsi et al., 2018).

Recently, a number of improvements have been made to EBM training techniques which have enabled EBMs to be trained on high-dimensional data. These include improvements to MCMC-based training (Nijkamp et al., 2019a), Score Matching (Song & Ermon, 2019), and noise-contrastive approaches (Gao et al., 2019). These improvements enabled new applications in domains such as protein structure prediction (Ingraham et al., 2019; Du et al., 2020), and provided benefit to fundamental problems in machine learning such as adversarial robustness, calibration, out-of-distribution detection (Grathwohl et al., 2019; Du & Mordatch, 2019), and semi-supervised learning (Song & Ou, 2018).

Despite this progress, we have little insight into the quality of the models we are learning. One solution is to indirectly measure quality of the model by evaluating performance on downstream discriminative tasks as advocated for in Theis et al. (2015), and recently applied to EBMs by Grathwohl et al. (2019). Another solution is to use metrics that rely on samples generated by expensive MCMC algorithms (Nijkamp et al., 2019a; Song & Ermon, 2019; Du & Mordatch, 2019).

In this work, we develop a unified approach to training and evaluating EBMs which addresses many of the aforementioned issues. We produce a measure of model fit that requires only an unnormalized model and data from the target distribution. This measure is given by a neural network which is trained to estimate the Stein Discrepancy between distributions. To our knowledge, our work is the first to empirically demonstrate that a neural network can be trained to reliably estimate the Stein Discrepancy in high dimensions. The resulting evaluation and training procedures outperform previous approaches at both tasks.

2. Problems with Sample-Based Training and Evaluation of Energy-based Models

EBMs and MCMC sampling have been closely intertwined since their inception. Much prior work trains EBMs by approximating the gradient of the model likelihood with:

$$\frac{\partial \log p_{\theta}(\mathbf{x})}{\partial \theta} = \mathbf{E}_{\mathbf{x}' \sim p_{\theta}(\mathbf{x})} \left[\frac{\partial E_{\theta}(\mathbf{x}')}{\partial \theta} \right] - \frac{\partial E_{\theta}(\mathbf{x})}{\partial \theta}, \quad (2)$$

where θ parameterizes the model, and the expectation on the right-hand-side of Equation 2 is estimated using MCMC. We refer to such approaches as Approximate Maximum Likelihood (AML). When the Markov chain is seeded from training data, this approach is referred to as Contrastive Divergence (CD). This gradient estimator has been used in

the past to train product-of-experts (Hinton, 2002) and Restricted Boltzmann Machines (Hinton et al., 2006). Recently, advances in generic gradient-based samplers (Welling & Teh, 2011) have allowed this gradient estimator to be used to train much more large-scale models (Du & Mordatch, 2019; Grathwohl et al., 2019; Ingraham et al., 2019).

Unfortunately, unless extreme care is taken (Jacob et al., 2017), estimating an expectation with a finite Markov chain will produce biased estimates. This leads to optimizing an objective other than the one we wish to. An MCMC sampler will only draw true samples from an unnormalized density if it is run for an infinite number of steps. When a finite number of steps are used, the distribution of the resulting samples can be arbitrarily far away from the true distribution parameterized by the model. This difference is a function of the model itself as well as the parameters of the sampler such as its step size, number of steps, and initialization distribution. This means the bias of this training objective cannot be easily quantified and is greatly complicated by the choice of sampler. This phenomenon has been explored in Nijkamp et al. (2019b) where the authors argue that training in this way actually learns an implicit sampler and not a density model. This is illustrated in Figure 1.

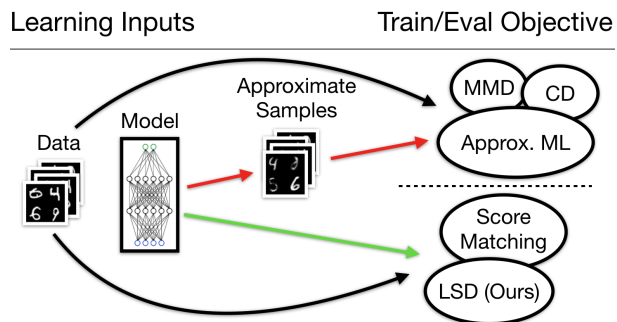


Figure 2. Cutting out the “middle-man” of approximate sampling can lead to simpler training and evaluation that is tied directly to the quality of our model and is not obfuscated by the parameters of an MCMC sampler.

3. Assessing Fit Without Samples or Normalizing Constants

To avoid the problems of sampler-based training and evaluation, we seek a measure of model fit that can be evaluated given a finite set of datapoints and an unnormalized density model. Surprisingly, such a measure exists and is based on the following, known as Stein’s Identity (Stein et al., 1972):

$$\mathbf{E}_{p(x)} [\nabla_x \log p(x)^T f(x) + \text{Tr}(\nabla_x f(x))] = 0 \quad (3)$$

where $f : \mathcal{R}^D \rightarrow \mathcal{R}^D$ is any function such that $\lim_{\|x\| \rightarrow \infty} p(x)f(x) = 0$. We refer to f as the critic. Note that $\nabla_x \log p(x)$ can be evaluated without computing the

normalizing constant of $p(x)$. If we replace the $p(x)$ inside of the expectation with a different distribution $q(x)$, then this expectation is zero for all f if and only if $p = q$.

Taking f to be the supremum over a class of functions \mathcal{F} , we can create a quantitative measure:

$$\mathbf{S}(p, q) = \sup_{f \in \mathcal{F}} \mathbf{E}_{p(x)} [\nabla_x \log q(x)^T f(x) + \text{Tr}(\nabla_x f(x))] , \quad (4)$$

which is known as the Stein Discrepancy (SD) (Gorham & Mackey, 2017).

If \mathcal{F} is taken to be a ball in a Reproducing Kernel Hilbert Space (RKHS), then a kernelized version of this discrepancy can be derived:

$$\begin{aligned} \mathbf{KSD}(p, q) = \mathbf{E}_{x, x' \sim p(x)} [& \nabla_x \log q(x)^T k(x, x') \nabla_{x'} \log q(x') \\ & + \nabla_x \log q(x)^T \nabla_{x'} k(x, x') \\ & + \nabla_x k(x, x')^T \nabla_{x'} \log q(x') \\ & + \text{Tr}(\nabla_{x, x'} k(x, x'))] , \quad (5) \end{aligned}$$

known as the Kernelized Stein Discrepancy (KSD), which has been used to build hypothesis tests to assess goodness-of-fit for unnormalized densities (Liu et al., 2016; Gorham & Mackey, 2017; Jitkrittum et al., 2017) and to learn implicit samplers for unnormalized densities (Hu et al., 2018).

4. Learning the Stein Discrepancy

The KSD allows us to circumvent the functional optimization in Equation 4 by using a kernel function $k(\cdot, \cdot)$. Unfortunately, it is known that methods based on distances and kernels quickly degrade in performance as dimension increases (Ramdas et al., 2015). Further, the power of tests based on these kernel methods is closely tied to their asymptotic run-time; the quadratic time test of Liu et al. (2016) significantly outperforms their linear-time variant, which prevents using the best-performing approach on large datasets.

Using the Stein Discrepancy directly has the potential to address both these issues. By employing a larger and more expressive class of functions \mathcal{F} , we can produce a more discriminative measure in high dimensions. In the functional form, we can estimate the discrepancy in linear time with respect to the number of examples.

We propose to parameterize the critic with a neural network f_ϕ and optimize its parameters to maximize

$$\mathbf{LSD}(f_\phi, p, q) = \mathbf{E}_{p(x)} [\nabla_x \log q(x)^T f_\phi(x) + \text{Tr}(\nabla_x f_\phi(x))] , \quad (6)$$

which we call the Learned Stein Discrepancy (LSD). We will then use this learned discrepancy to evaluate and train unnormalized models.

Choosing \mathcal{F} To estimate a Stein Discrepancy as in Equation 4, we must optimize the critic over a bounded space of functions \mathcal{F} . In general, unconstrained neural networks do not fall into this category, so additional care must be taken. This can be accomplished in many ways such as with weight-clipping, or spectral normalization (Miyato et al., 2018). In this work, we optimize critic networks within

$$\mathcal{F} = \{f : \mathbf{E}_{p(x)} [f(x)^T f(x)] < \infty\} , \quad (7)$$

the space of functions whose squared norm has finite expectation under the data distribution. This constraint will be enforced by placing an \mathcal{L}_2 regularizer on our critic’s output

$$\mathcal{R}_\lambda(f_\phi) = \lambda \mathbf{E}_{p(x)} [f_\phi(x)^T f_\phi(x)] \quad (8)$$

with strength λ . Thus, our critic will be trained to maximize $\mathbf{LSD}(f_\phi, p, q) - \mathcal{R}_\lambda(f_\phi)$.

Under this class of functions, Hu et al. (2018) prove the optimal critic takes the following form:

$$f_\phi(x) = \frac{1}{2\lambda} (\nabla_x \log q(x) - \nabla_x \log p(x)) . \quad (9)$$

Thus, the optimal critic depends on λ only up to a constant multiplier. While this theory is convenient, it does not tell us if we can learn this function given finite data and compute – requirements for this to be a useful discrepancy for model evaluation and training. In this work, we demonstrate empirically that we can learn this function from finite data on a wide variety of datasets and models. In Figure 3 we show that the theoretically optimal critic can be recovered while estimating the Stein Discrepancy between 100 dimensional Gaussian distributions.

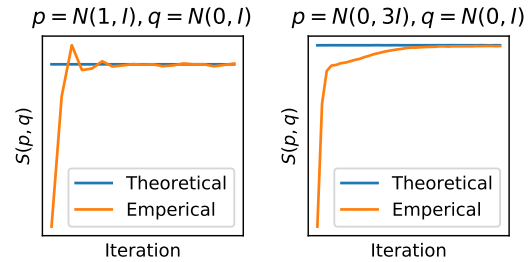


Figure 3. Training a neural net to estimate $\mathbf{S}(p, q)$ on 100-dimensional data. In both cases, a near-optimal critic is learned and the true discrepancy is closely approximated.

We achieve similar results on more complicated distributions such as RBMs and Normalizing Flows trained on image data. See Sections 7.2 and 7.3 for details.

Efficient Estimation The $\text{Tr}(\nabla_x f(x))$ term in Eq. 6 is expensive to compute, requiring $O(D)$ vector-Jacobian products (or “backward-passes”), where D is the dimension of

the data. Fortunately, an efficient, unbiased estimator exists, known as Hutchinson’s estimator (Hutchinson, 1990). This estimator has been widely used in the machine learning community in recent years (Grathwohl et al., 2018; Tsitsulin et al., 2019; Han et al., 2017). The estimator, which requires only one vector-Jacobian product to compute, is the single-sample Monte-Carlo estimator derived from the following identity:

$$\text{Tr}(\nabla_x f(x)) = \mathbf{E}_{N(\epsilon|0,1)} [\epsilon^T \nabla_x f(x) \epsilon] \quad (10)$$

which can be computed efficiently since $\epsilon^T \nabla_x f(x)$ is a vector-Jacobian product. Thus, during the critic-learning phase, we replace the LSD objective (Equation 6) with

$$\text{LSDE}(f_\phi, p, q) = \mathbf{E}_{p(x)N(\epsilon|0,1)} [\nabla_x \log q(x)^T f_\phi(x) + \epsilon^T \nabla_x f_\phi(x) \epsilon] \quad (11)$$

an Efficient LSD version which can be estimated and optimized on mini-batches of data sampled from $p(x)$ for the same asymptotic time cost as evaluating f and $\nabla_x \log q(x)$ once on each sample in the mini-batch.

5. Model Evaluation with LSD

Now that we have a procedure for estimating the Stein Discrepancy, we provide two applications related to model evaluation: comparing the performance of models on held-out evaluation data, and goodness-of-fit testing.

5.1. Model Comparison

In this setting we are given two or more unnormalized models $\{q_i\}_{i=1}^M$, and a finite set of samples $\{x_i\}_{i=1}^n$ from the target distribution $p(x)$ which $\{q_i\}_{i=1}^M$ are trying to approximate. The goal is to estimate which q_i approximates the data samples better.

We phrase this evaluation task as a standard learning problem. We split the data into training, validation, and testing sets, and train a critic f_{ϕ_i} on model q_i , using the validation data to do model selection. We then compute the LSD on the test data for each model and score the models by this value.

Due to the nature of our training objective, the variance of the LSD must be taken into consideration when monitoring for over-fitting. We propose an uncertainty-aware model selection procedure where we choose the model which maximizes $\mu - \sigma$ where μ and σ are the sample mean and standard deviation of the LSD on the validation data. We provide a more in-depth discussion in Appendix A.

Our approach is summarized in Algorithm 1. See Sections 7.2, 7.3 for experimental result for experimental results, where we find that LSD can scale to complicated models of high-dimensional data such as images.

Algorithm 1 LSD Model Comparison

Input: Critic architecture f_ϕ , models $\{q_i\}_{i=1}^M$, data $\mathbf{x} = \{x_i\}_{i=1}^n$, \mathcal{L}_2 regularization hyperparameter λ
Output: Estimated Stein Discrepancies, $\{\mathbf{S}(p, q_i)\}_{i=1}^M$
 Split \mathbf{x} into $\mathbf{x}_{\text{train}}$, \mathbf{x}_{val} , and \mathbf{x}_{test}
for model q_i in $\{q_i\}_{i=1}^M$ **do**
 find $\phi_i = \text{argmax}_\phi \text{LSDE}(f_\phi, \mathbf{x}_{\text{train}}, q_i) - \mathcal{R}_\lambda(f_\phi)$
 (using \mathbf{x}_{val} for model selection)
 $s_i = \text{LSD}(f_{\phi_i}, \mathbf{x}_{\text{test}}, q_i)$
end for
 Rank models in increasing order of s_i .

5.2. Goodness-Of-Fit Testing

Given an unnormalized model $q(x)$ and a finite set of samples $\{x_i\}_{i=1}^n$ from an unknown distribution $p(x)$, the problem of Goodness-Of-Fit (GoF) testing asks us to decide between two hypotheses:

$$H_0 : p = q, \quad H_1 : p \neq q.$$

GoF is a standard problem in statistics. An ideal hypothesis test will reject H_0 whenever $p \neq q$, even if p and q are quite similar. In settings relevant to the machine learning community, GoF becomes particularly challenging as we often deal with very complicated, high-dimensional data.

We use LSD to develop a procedure for solving GoF problems. Assuming we are given a critic f , we define $s_f^q(x) = \nabla_x \log q(x)^T f(x) + \text{Tr}(\nabla_x f(x))$. Following Equation 3, this problem reverts to a test of

$$H_0 : \mathbf{E}_{p(x)} [s_f^q(x)] = 0 \quad H_1 : \mathbf{E}_{p(x)} [s_f^q(x)] \neq 0$$

which is a simple one-sample location test. This test is carried out by computing the statistic $t = \sqrt{n} \frac{\mu_s}{\sigma_s}$ where μ_s and σ_s are the sample mean and standard deviation of $s_f(x)$ evaluated over the dataset. For sufficiently large n , we have $t \sim N(0, 1)$ under H_0 , thus for a given test confidence α we should reject H_0 if $t < \Phi(1 - \alpha)$ where Φ is the inverse CDF of the standard Normal distribution.

As above, to obtain the critic, we parameterize it as a neural network f_ϕ and train its parameters on a subset of the given data. In Section 5.1 we were interested in estimating $\mathbf{S}(p, q)$ directly. Here we are only interested in correctly choosing H_0 or H_1 , so we should optimize f_ϕ to minimize the probability we accept H_0 given H_1 is true. The inverse of this quantity is known as the *test power*. Given that the test statistic t is asymptotically normal under both H_1 and H_0 we can follow the same argument as Sutherland et al. (2016); Jitkrittum et al. (2017), and maximize the test power by optimizing

$$\mathcal{P}(f_\phi, p, q) = \frac{\mathbf{E}_{p(x)} [s_{f_\phi}^q(x)]}{\sigma_{p(x)} [s_{f_\phi}^q(x)]}. \quad (12)$$

We split our finite sample set into disjoint sets $\mathbf{x}_{\text{train}}$, \mathbf{x}_{val} , and \mathbf{x}_{test} , using $\mathbf{x}_{\text{train}}$ for training and \mathbf{x}_{val} for model selection. Then, given our learned f_ϕ we run the test on x_{test} . A summary of our test can be seen in Algorithm 2. As before, we use Hutchinson’s estimator for the trace term in $s_f^q(x)$. Because this increases the variance of $s_f^q(x)$, optimizing the objective maximizes a lower-bound on the test power.

Experimental details and results can be found in Section 7.1.

Algorithm 2 LSD Goodness of Fit Test

Input: Critic architecture f_ϕ , model q , $\mathbf{x} = \{x_i\}_{i=1}^n$, \mathcal{L}_2 regularization hyperparameter λ , confidence α

Output: Decision whether $p = q$

Split \mathbf{x} into $\mathbf{x}_{\text{train}}$, \mathbf{x}_{val} , and \mathbf{x}_{test}

maximize $\phi = \operatorname{argmax}_\phi \mathcal{P}(f_\phi, \mathbf{x}_{\text{train}}, q)$ (using \mathbf{x}_{val} for model selection)

Compute $t = \sqrt{n} \frac{\mathbf{E}[s(\mathbf{x}_{\text{test}})]}{\sigma[s(\mathbf{x}_{\text{test}})]}$

If $t > \Phi(1 - \alpha)$ reject H_0 , else accept H_0

6. Training Unnormalized Models with LSD

Above we have developed a method which can effectively quantify the quality of fit of an unnormalized model to a fixed set of datapoints. We now extend this method to develop a way of learning the parameters of an unnormalized model to best fit the data. Unlike sampling based approaches, we produce an objective that is 0 in expectation if and only if the data and model distributions are equivalent, given our critic is suitably expressive. Ideally, we will minimize

$$\mathcal{L}(\theta) = \sup_\phi \mathbf{LSD}(f_\phi, p, q_\theta) - \mathcal{R}_\lambda(f_\phi) \quad (13)$$

with respect to θ but this is infeasible due to the supremum on the right-hand-side. Instead, we follow recent work on minimax optimization (Goodfellow et al., 2014) and iteratively update the model $q_\theta(x) = \exp(-E_\theta(x))/Z$ and the critic function f_ϕ in an alternating fashion where the critic is trained to more accurately estimate the Stein Discrepancy and the model is trained to minimize the critic’s estimate of the discrepancy. The proposed training algorithm is summarized in Algorithm 3.

Our approach can effectively train unnormalized models on complicated high-dimensional datasets without requiring samples of the model – only needing access to the model’s score function. In contrast, score matching (which also does not require sampling), requires the Hessian of $\log q_\theta$ for training. For many distributions of interest this computation can be numerically unstable leading to training issues (see Section 8.1 for more details). Further, optimizing functions implicitly through their higher-order derivatives can be challenging as they are often sparse or discontinuous. See Figure 4 for examples of toy densities trained with LSD.




Figure 4. Density models trained using LSD. Top: Data. Bottom: Learned densities.

Algorithm 3 LSD Training

Input: Critic architecture f_ϕ , model q_θ , data $\mathbf{x} = \{x_i\}_{i=1}^n \sim p(x)$, \mathcal{L}_2 regularization hyperparameter λ , training iterations T , critic training iterations C

Output: Parameters θ such that $q_\theta \approx p$

for T iterations **do**

for C iterations **do**

 Sample mini-batch \mathbf{x}'

 Update ϕ with $\nabla_\phi (\mathbf{LSDE}(f_\phi, \mathbf{x}', q_\theta) - \mathcal{R}_\lambda(f_\phi))$

end for

 Sample mini-batch \mathbf{x}'

 Update θ with $-\nabla_\theta \mathbf{LSD}(f_\phi, \mathbf{x}', q_\theta)$

end for

 Return resulting model q_θ

7. Model Evaluation Experiments

We run a number of experiments to demonstrate the utility of the LSD. First, we demonstrate that the LSD can be used to build a competitive test for goodness-of-fit which scales more favorably to high-dimensional data than prior linear-time methods. Next we show that LSD can be used to provide a fine-grained model evaluation tool which is sensitive to small differences in high-dimensional models.

Unless otherwise stated, all critics f_ϕ are MLPs with 2 hidden layers and 300 units per layer. We use the Swish (Ramachandran et al., 2017) nonlinearity throughout.

7.1. Hypothesis Testing

We compare our linear-time hypothesis testing method from Section 5.2 with a number of kernel-stein approaches; the quadratic-time Kernelized Stein Discrepancy (KSD), its linear-time variant (LKSD) (Liu et al., 2016), and the linear-time Finite Set Stein Discrepancy (FSSD) (Jitkrittum et al., 2017). We also compare against a kernel-MMD test (Gretton et al., 2012) which requires MCMC sampling from the




Figure 5. Hypothesis testing results. Test confidence 0.05. Perturbed RBMs of increasing data dimension. Perturbation magnitude on the x -axis, rejection rate on the y -axis. Number of datapoints $n = 1000$. Ideal behavior is a 5% rejection when perturbation is 0 and close to 100% rejection otherwise. In high dimensions our linear-time **LSD** matches the performance of the quadratic-time **KSD**.

model. We test these approaches in their ability to determine whether or not a set of samples was drawn from a given Gaussian-Bernoulli Restricted Boltzmann Machine (Cho et al., 2013). This is an unnormalized latent-variable model

$$p(x, h) = \frac{1}{Z} \exp\left(\frac{1}{2}x^T B h + b^T x + c^T h - \frac{1}{2}\|x\|^2\right)$$

$$\nabla_x \log p(x) = b - x + B \cdot \tanh(B^T x + c)$$

with parameters B, b, c , whose gradients can be efficiently computed making it an ideal candidate for evaluating methods such as ours. We randomly sample the parameters of the model, then draw a set of $n = 1000$ samples. We then perturb the weights of the model with Gaussian noise of standard deviation in $[0, 0.01, 0.02, 0.04, 0.06]$ and our tests must determine if the samples were drawn from this model. We perform this test with RBMs of increasing visible dimension x and hidden dimension h from $x, h = \{(50, 40), (100, 80), (200, 100)\}$.

As can be seen in Figure 5, our proposed hypothesis test performs comparably to the linear-time kernel methods at 50 dimensions and begins to dominate those approaches as dimensionality is increased, matching the performance of the quadratic-time test (we note that this quadratic-time test takes much longer to run than our linear-time method). We also see that while the MMD test performs comparably at 50 dimensions, as we increase dimensionality the performance quickly drops below all of the Stein approaches. This result further supports our claim that MCMC sampling should not be relied upon in high dimensional settings.

We run further experiments to empirically verify the normality of our test statistic under H_0 . These results can be found in Appendix C.

7.2. RBM Evaluation

We now demonstrate **LSD**'s ability to rank and evaluate the fit of unnormalized models on fixed test data. Again, we experiment with Gaussian-Bernoulli RBMs. As above we randomly initialize an RBM, draw $n = 1000$ samples,

perturb its weights with increasing Gaussian perturbations, and then and provide a score for each model. This score should increase as the perturbation becomes larger, starting at 0 when $p = q$. As above, we experiment with RBMs of dimension 50, 100, and 200.

We compare **LSD** (which approximates the Stein Discrepancy over \mathcal{L}_2 functions) with a linear-time and a quadratic-time estimate of the KSD using RBF Kernels with learned bandwidth. We also compare with the theoretical upper-bound on the **LSD** from Equation 9.

We test all approaches using $n = 1000$ samples as this was the largest n that could feasibly be used for the quadratic-time KSD. In all dimensions, we find that **LSD** is a much stronger measure of model quality. It is able to closely match the theoretical upper bound even in higher dimensions. Further, it is much more discriminative for small-scale perturbations than the kernel methods. Results can be seen in Figure 6. We plot the reported scores with error-bars indicating standard deviation of the sample mean $\frac{\sigma}{\sqrt{n_{\text{test}}}}$.

7.3. Normalizing Flow Evaluation

For the **LSD** to be an effective discrepancy for use on problems of scale, it must be able to distinguish between relatively similar, high-dimensional models. To assess **LSD**'s ability to do so, we trained a number of normalizing flow models based on the Glow architecture (Kingma & Dhariwal, 2018). We save these models throughout training and record their log-likelihoods on the MNIST test dataset. We compare the standard likelihood evaluation with the **LSD** evaluation procedure. We split the MNIST test set into partitions to train, validate, and test the **LSD** critic. This ensures that the scores from **LSD** were given access to the exact same data as the likelihood evaluation procedure. We also compare with the objective of Score Matching (Hyvärinen, 2005) to provide a baseline method which does not require the model's normalizing constant.

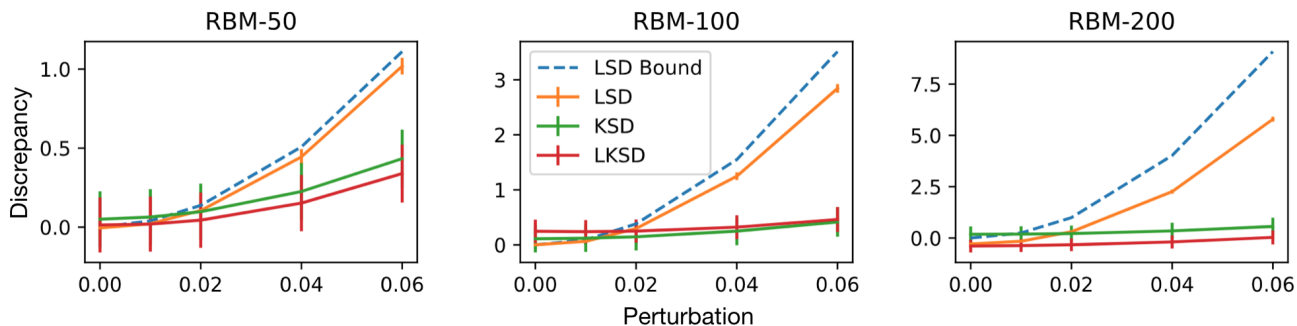


Figure 6. Model evaluation results for Gaussian-Bernoulli RBMs. Predicted discrepancy should increase as perturbation increases. **LSD** reliably increases, approaching the theoretical ground-truth value. **LSD** provides much more certain results than both linear-time and quadratic-time kernel-based approaches.

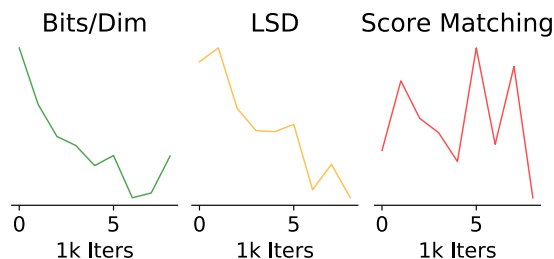


Figure 7. Flow model evaluation. Left to right: Bits/Dim, LSD, Score Matching. Y-axis is removed because scales are not comparable.

Results can be seen in Figure 7. We find that the score reported by **LSD** goes down throughout training, tracking the negative log-likelihood. Conversely, score matching reports a score which is not highly correlated with likelihood, demonstrating that **LSD** can be a more effective metric to compare unnormalized models than previous approaches.

8. Model Training Experiments

Here we demonstrate that minimizing **LSD** is an effective method for training unnormalized models which scales to high dimensional data more effectively than previously proposed approaches. We have trained deep EBMs on some simple toy 2D densities using **LSD** which can be seen in Figure 4. Below we present quantitative results on some more challenging problems.

8.1. Linear ICA

The linear Independent Components Analysis (ICA) model is commonly used to quantitatively evaluate the performance of methods for training unnormalized models (Gutmann & Hyvärinen, 2010; Hyvärinen, 2005; Ceylan & Gutmann, 2018). It consists of a simple generative process

$$z \sim \text{Laplace}(0, 1), \quad x = Wz$$

where the model parameter W is a $D \times D$, non-singular mixing matrix. The log-density of a datapoint x under this model is

$$\log p(x; W) = \log p_z(W^{-1}x) - \log |W| \quad (14)$$

where $p_z(\cdot)$ is the PDF of the Laplace(0, 1) distribution.

We train linear ICA models on randomly sampled mixing matrices using various methods and compare their performance as dimension D is increased from 5 to 50. To provide an upper-bound on possible performance, we train with brute-force Maximum Likelihood (ML) (which requires inverting the parameter matrix at each iteration). We compare **LSD** with other approaches to train unnormalized models; Noise-Contrastive Estimation (NCE) (Gutmann & Hyvärinen, 2010), Conditional Noise-Contrastive Estimation (CNCE) (Ceylan & Gutmann, 2018), Score Matching (SM) (Hyvärinen, 2005).

Results can be seen in Figure 8. We find that for smaller D these methods behave comparably, and all arrive at a similar solution to maximum likelihood. For D above 20, other unnormalized training methods fail to achieve the same level of performance as ML while **LSD** does. For D above 30, SM quickly diverged due to instabilities that arose with computing the second derivatives of the Laplace log-density. Our approach only requires access to the first derivatives of the model and the critic, avoiding these issues. Further experimental details can be found in Appendix B.5.

8.2. Preliminary Image Modeling Results

We now show that **LSD** can be used to train much more expressive unnormalized models. We train deep energy-based models with **LSD** on MNIST and FashionMNIST. The models take the form of $q_\theta(x) = \exp(-E_\theta(x))/Z$, where the energy-function $E_\theta(x)$ is parameterized by a neural network with a single output variable. We parameterize the critic f_ϕ with a neural network which maps $R^D \rightarrow R^D$.

After training, we draw approximate samples with a tem-




Figure 8. Linear ICA training results. Left: Table comparing test set log-likelihoods of linear ICA models with various training methods. **LSD** closely tracks the performance maximum likelihood training while other unnormalized methods fall behind or diverge. Right: Learning curves for 50-dimensional ICA. While slower to converge than maximum likelihood, **LSD** cleanly converges to the same result.




Figure 9. Samples from deep EBMs trained with LSD. Left: MNIST, Right: FashionMNIST.

pered SGLD sampler. These can be seen in Figure 9. Without the use of a sampler, **LSD** has trained a model which captures the various modes of the data distribution and can be used to draw compelling samples. Of course, these MCMC samples are susceptible to all of the issues we have mentioned previously in this work. Their quality is tightly coupled with the post-hoc sampler’s parameters and we make no claims that these represent true samples from our models.

While these samples are not state-of-the-art, they do showcase the ability of **LSD** to scale to high-dimensional problems. Scaling further to large natural images will potentially require tricks from the GAN literature, the use of convolutional architectures, and more advanced samplers. We leave this for future work. We refer the reader to Appendix B.6 for experimental details, and Appendix D.1 for further samples and experiments using **LSD** to train RBMs on image datasets.

9. Related Work

Sampler-Free EBM Training Ours is not the first method to train unnormalized densities without MCMC sampling. Score matching (SM) (Hyvärinen, 2005) matches the gradients of the model density to the data density, circumventing the computation of the normalizing constant.

Unlike our approach, SM requires the computation of the unnormalized density’s Hessian trace, which can be expensive and unstable (see our ICA results in Section 8.1). An alternative interpretation of SM also exists known as Denoising Score Matching (Vincent, 2011) which approximates the SM objective without the Hessian term. This has been applied to deep EBMs (Saremi et al., 2018) and recently scaled to high resolution image data (Song & Ermon, 2019; Li et al., 2019) with compelling results.

These methods attempt to minimize the Fisher Divergence $F(p, q) = E_{p(x)}[|\|\nabla_x \log p(x) - \nabla_x \log q(x)\|_2^2]$, which can be viewed a special case of the Stein Discrepancy with a specific, fixed critic function (Liu et al., 2016). Barp et al. (2019) showed that other EBM training methods (such as approximate maximum likelihood) can be viewed as minimizing a Stein Discrepancy with respect to a different class of critics.

Approximating Stein Discrepancies A second highly relevant prior contribution is the Stein Neural Sampler (Hu et al., 2018). which presents a setup very similar to our own but with the opposite motivation. In their work, one is given an unnormalized density $q(x)$ (such as the posterior of a Bayesian model) and seeks to learn an implicit sampler $x = g_\theta(z)$ such that $x \sim q(x)$. This is achieved through a training procedure like ours presented in Section 6 but with $q(x)$ fixed, and the gradients of **LSD** back-propagated back through the samples, to the parameters of the sampler.

Our work benefited significantly from the theoretical results of their work, but their method was not particularly scalable due to the brute-force evaluation of a Jacobian trace. We solve this issue with Hutchinson’s estimator (Eq. 11). We are particularly excited about this line of work and believe there is much promise to further scale these sampling methods in this way.

Variational Inference Computing and minimizing Stein Discrepancies has been studied in the context of variational

inference for both sampling (Liu & Wang, 2016) and amortized inference. In particular, Ranganath et al. (2016) introduce an objective for inference similar to LSD but utilize a different class of critic function and not dot apply many of the techniques we used for scalability. We expect using our class of critics and these techniques could improve the performance of their approach.

GANs and Adversarial Optimization Generative adversarial networks (GANs) (Goodfellow et al., 2014) (and specifically, their extension Wasserstein GANs (Arjovsky et al., 2017)) have many similarities to LSD. They both train a model by minimizing a discrepancy between the model distribution and the data distribution which is defined through a critic function. Both approaches parameterize this critic with a neural network and train this critic online with the model in an alternating fashion. The main difference lies in the properties of the model; the WGAN trains an implicit sampler while LSD trains an unnormalized density model.

GANs have successfully scaled to very high dimensional datasets (Brock et al., 2018) and their success has motivated considerable progress in minimax optimization. Due to the similarity of the approaches, we are excited about using these techniques to scale LSD to more challenging problems than we have attacked in this work.

10. Conclusion

In this work we have presented LSD, a novel and scalable way to train and evaluate unnormalized density models without the need for sampling or kernel-selection heuristics. LSD outperforms state-of-the-art methods for high-dimensional hypothesis testing. We show that LSD tracks likelihood much better than common objective functions used to train unnormalized density models. Finally, we show that LSD can be used to train unnormalized density models efficiently in high dimensions.

11. Acknowledgements

We thank Renjie Liao and Murat Erdogdu for useful discussion. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute (www.vectorinstitute.ai/#partners).

References

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

Barp, A., Briol, F.-X., Duncan, A., Girolami, M., and Mackey, L. Minimum stein discrepancy estimators. In

Advances in Neural Information Processing Systems, pp. 12964–12976, 2019.

- Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Ceylan, C. and Gutmann, M. U. Conditional noise-contrastive estimation of unnormalised models. *arXiv preprint arXiv:1806.03664*, 2018.
- Cho, K. H., Raiko, T., and Ilin, A. Gaussian-bernoulli deep boltzmann machine. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–7. IEEE, 2013.
- Du, Y. and Mordatch, I. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.
- Du, Y., Meier, J., Ma, J., Fergus, R., and Rives, A. Energy-based models for atomic-resolution protein conformations. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=S1e_9xrFvS.
- Falorsi, L., de Haan, P., Davidson, T. R., De Cao, N., Weiler, M., Forré, P., and Cohen, T. S. Explorations in homeomorphic variational auto-encoding. *arXiv preprint arXiv:1807.04689*, 2018.
- Gao, R., Nijkamp, E., Kingma, D. P., Xu, Z., Dai, A. M., and Wu, Y. N. Flow contrastive estimation of energy-based models. *arXiv preprint arXiv:1912.00589*, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gorham, J. and Mackey, L. Measuring sample quality with kernels. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1292–1301. JMLR. org, 2017.
- Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- Grathwohl, W., Wang, K.-C., Jacobsen, J.-H., Duvenaud, D., Norouzi, M., and Swersky, K. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263*, 2019.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.

- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 297–304, 2010.
- Han, I., Malioutov, D., Avron, H., and Shin, J. Approximating spectral sums of large-scale matrices using stochastic chebyshev approximations. *SIAM Journal on Scientific Computing*, 39(4):A1558–A1585, 2017.
- Hinton, G. E. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.
- Hinton, G. E., Osindero, S., and Teh, Y.-W. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- Hu, T., Chen, Z., Sun, H., Bai, J., Ye, M., and Cheng, G. Stein neural sampler. *arXiv preprint arXiv:1810.03545*, 2018.
- Hutchinson, M. F. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 19(2):433–450, 1990.
- Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.
- Ingraham, J., Riesselman, A., Sander, C., and Marks, D. Learning protein structure with a differentiable simulator. In *International Conference on Learning Representations*, 2019.
- Jacob, P. E., O’Leary, J., and Atchadé, Y. F. Unbiased markov chain monte carlo with couplings. *arXiv preprint arXiv:1708.03625*, 2017.
- Jitkrittum, W., Xu, W., Szabó, Z., Fukumizu, K., and Gretton, A. A linear-time kernel goodness-of-fit test. In *Advances in Neural Information Processing Systems*, pp. 262–271, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pp. 10215–10224, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- LeCun, Y., Chopra, S., Ranzato, M., and Huang, F.-J. Energy-based models in document recognition and computer vision. In *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, volume 1, pp. 337–341. IEEE, 2007.
- Li, Z., Chen, Y., and Sommer, F. T. Annealed denoising score matching: Learning energy-based models in high-dimensional spaces. *arXiv preprint arXiv:1910.07762*, 2019.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in neural information processing systems*, pp. 2378–2386, 2016.
- Liu, Q., Lee, J., and Jordan, M. A kernelized stein discrepancy for goodness-of-fit tests. In *International conference on machine learning*, pp. 276–284, 2016.
- Lucas, J., Tucker, G., Grosse, R., and Norouzi, M. Understanding posterior collapse in generative latent variable models. 2019.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Nijkamp, E., Hill, M., Han, T., Zhu, S.-C., and Wu, Y. N. On the anatomy of mcmc-based maximum likelihood learning of energy-based models. *arXiv preprint arXiv:1903.12370*, 2019a.
- Nijkamp, E., Zhu, S.-C., and Wu, Y. N. On learning non-convergent short-run mcmc toward energy-based model. *arXiv preprint arXiv:1904.09770*, 2019b.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- Osadchy, M., Cun, Y. L., and Miller, M. L. Synergistic face detection and pose estimation with energy-based models. *Journal of Machine Learning Research*, 8(May): 1197–1215, 2007.
- Ramachandran, P., Zoph, B., and Le, Q. V. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- Ramdas, A., Reddi, S. J., Póczos, B., Singh, A., and Wasserman, L. On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Ranganath, R., Tran, D., Altsosaar, J., and Blei, D. Operator variational inference. In *Advances in Neural Information Processing Systems*, pp. 496–504, 2016.

- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Saremi, S., Mehrjou, A., Schölkopf, B., and Hyvärinen, A. Deep energy estimator networks. *arXiv preprint arXiv:1805.08306*, 2018.
- Scellier, B. and Bengio, Y. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. *arXiv preprint arXiv:1907.05600*, 2019.
- Song, Y. and Ou, Z. Learning neural random fields with inclusive auxiliary generators. *arXiv preprint arXiv:1806.00271*, 2018.
- Song, Y., Garg, S., Shi, J., and Ermon, S. Sliced score matching: A scalable approach to density and score estimation. *arXiv preprint arXiv:1905.07088*, 2019.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Stein, C. et al. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*. The Regents of the University of California, 1972.
- Sutherland, D. J., Tung, H.-Y., Strathmann, H., De, S., Ramdas, A., Smola, A., and Gretton, A. Generative models and model criticism via optimized maximum mean discrepancy. *arXiv preprint arXiv:1611.04488*, 2016.
- Theis, L., Oord, A. v. d., and Bethge, M. A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844*, 2015.
- Tsitsulin, A., Munkhoeva, M., Mottin, D., Karras, P., Bronstein, A., Oseledets, I., and Müller, E. The shape of data: Intrinsic distance for data distributions. *arXiv preprint arXiv:1905.11141*, 2019.
- Vincent, P. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine Learning (ICML)*, pp. 681–688, 2011.
- Xie, J., Lu, Y., Zhu, S.-C., and Wu, Y. A theory of generative convnet. In *International Conference on Machine Learning*, pp. 2635–2644, 2016.
- Xie, J., Zheng, Z., Gao, R., Wang, W., Zhu, S.-C., and Nian Wu, Y. Learning descriptor networks for 3d shape synthesis and analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8629–8638, 2018.
- Xie, J., Zhu, S.-C., and Wu, Y. N. Learning energy-based spatial-temporal generative convnets for dynamic patterns. *IEEE transactions on pattern analysis and machine intelligence*, 2019.

A. Model Selection

Here we elaborate on some considerations that must be taken when training the **LSD** on finite data. Like any over-parameterized machine learning model, our critic networks are susceptible to over-fitting.

Our training objective repeated is:

$$\text{LSD}(f_\phi, q, p) = \mathbf{E}_{p(x)}[\nabla_x \log q(x)^T f_\phi(x) + \text{Tr}(\nabla_x f_\phi(x))].$$

If the regularization parameter λ is not set large enough, then this objective can be increased on the finite training data by simply adding a scalar multiplier onto the output of f_ϕ . This has the effect of increasing the variance of the **LSD**. Since we would like to select the model which maximizes the expectation of the **LSD**, a natural validation statistic would simply be the mean of the **LSD** on the validation set. Unfortunately, as the variance of the **LSD** increases, so does the variance of this statistic making this statistic an increasingly unreliable predictor of the model’s true performance. This means it is possible for this validation statistic to continue to increase as the model is over-fitting.

To combat this behavior, we propose a variance-aware validation statistic which is simply $\mu - \sigma$ or the mean of the **LSD** minus its standard deviation computed on the validation data. This statistic will decrease as the estimator variance increases, solving the issue.

We demonstrate this visually in Figure 10. We train a critic to estimate the **LSD** between data sampled from a 10-dimensional $N(0, 1)$ and a $N(1, 1)$ model q . The training dataset consists of 100 examples, so a flexible neural network critic will certainly over-fit. We see that as the model trains, the **LSD** evaluated on the 100-example training set quickly passes the theoretical bound (which is the true Stein Discrepancy), indicating over-fitting. The variance of the **LSD** validation data increases with over-fitting. We see that the best model is selected by our variance-aware model selection procedure.

A.1. Hypothesis Testing

We note that the above model selection procedure is *not required* when training critics for hypothesis testing. The test power objective (Equation 12) explicitly minimizes its variance so standard model selection may be used.

B. Experimental Details

B.1. Hypothesis Testing

We initialize RBMs following Liu et al. (2016); Jitkrittum et al. (2017). We select the entries of B uniformly from $\{+1, -1\}$ and we sample b and c from a standard Gaussian distribution. To test in a setting where the quadratic time

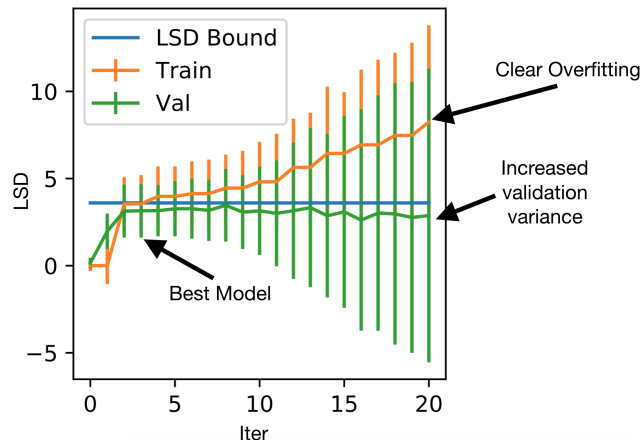


Figure 10. Visual motivation of our variance-aware model selection procedure. Plots of mean and standard deviation of the **LSD** computed on training and validation data.

test can be used for comparison, we use $n = 1000$ samples for our test.

For the adaptive kernel methods, we use 200 of these samples for learning the kernel parameters and use the remaining 800 samples to compute the test statistic. Given that **LSD** has many more parameters to fit, we use 800 samples for training, 100 samples for validation and 100 samples for running the final test.

We parameterize the critic f_ϕ as a 2-layer MLP with 300 units per layer and use the Swish nonlinearity. Due to the small amount of training data available, we regularize the model with dropout (Srivastava et al., 2014) and weight decay of strength .0005. We train using the Adam optimizer (Kingma & Ba, 2014) for 1,000 iterations. We do not use mini-batches, each iteration uses the entire training set. Every 100 iterations, we compute $\mathcal{P}_\phi(\mathbf{x}_{\text{val}})$ (Equation 12) on the validation data and select the critic with the highest value. We then compute the test statistic using this critic. For all tests we let $\lambda = .5$

Rejection rates presented are the average over 200 random tests.

B.2. RBM Evaluation

Experimental setup, (train/validation/test) splits, and models are the same as above but the objectives for training and testing differ. The hypothesis testing models (**LSD** and baselines) are trained to optimize test power which is the ratio of the mean of the estimated discrepancy over its standard deviation. This is the optimal objective when we wish to make a binary choice. Here we are interested in ranking models. Thus we train the baselines to maximize the dis-

crepancy they report. This means maximizing Equation 5 with respect to the kernel parameters. For the **LSD** models, we maximize Equation 11. For all models we let $\lambda = .5$.

We must be aware of estimator variance when doing model selection. We follow the procedure outlined in Appendix A to do model selection using our validation set. Over-fitting is not an issue with the kernel methods, so this is not done for the baselines.

B.3. Flow Evaluation

We train a small normalizing flow models based on the Glow architecture (Kingma & Dhariwal, 2018). The flow is an additive multi-scale architecture, uses reverse shuffle dimension splitting, actnorm, has 3 stages with 4 blocks per stage. Each block has 3 convolutional layers with a kernel size of: $(3 \times 3, 1 \times 1, 3 \times 3)$ respectively and 128 channels. The flow was trained with the Adam optimizer, a learning rate of .0001. We checkpoint the model 10 times throughout training, once every 1000 iteration. The models get between 2.0 and 1.8 bits/dim.

We take the 10,000 examples in the MNIST test set and split them into subsets for training, validation, and final score estimation. We use 8000, 1000, and 1000 examples for this, respectively. The critic was a 2-layer MLP with 1000 units per layer and used the Swish nonlinearity. We train for 100 passes through the data using a batch size of 128, validating after every pass.

B.4. Toy Densities

The toy densities in Figure 4 and the critic used to train them were parameterized by a 2-layer MLP with 300 units and Swish nonlinearity. The models were trained with the Adam optimizer (Kingma & Ba, 2014) for 10,000 iterations with a step size of .001. In both the critic and model, we set the Adam momentum parameters as $\beta_1 = .5, \beta_2 = .9$. The critic was updated 5 times for every model update. To ensure that the our model density can be normalized, we multiply the density we parameterize by a Gaussian distribution’s density:

$$q_\theta(x) = \frac{\exp(-E_\theta(x))N(x; \mu, \sigma^2)}{Z} \quad (15)$$

and we also learn the parameters μ and σ^2 with our model. We find empirically this Gaussian learns to cover the data and the energy function learns to push down the Gaussian’s likelihood at areas away from the data.

B.5. Linear ICA

ICA weight matrices were generated as random Gaussian matrices such that their condition number was smaller than the dimension of the matrix. Results shown are mean over

5 random seeds.

To make score matching efficient, we also utilize Hutchinson’s estimator when computing the Hessian trace, thus the objective we use to train is

$$\begin{aligned} \mathcal{J}_\theta(x) &= \frac{1}{2} \|\nabla_x \log q_\theta(x)\|^2 + \epsilon^T \mathcal{H}_x(\log q_\theta(x)) \epsilon^T \\ \epsilon &\sim N(0, I). \end{aligned} \quad (16)$$

This modified score matching objective has been studied previously and is known as Sliced Score Matching (Song et al., 2019).

All baselines were optimized for test performance. For all methods, the learning rate was selected from [.001, .0001, .00001].

For Noise Contrastive Estimation (Gutmann & Hyvärinen, 2010), we fit a Gaussian to the training data and use this as our noise distribution.

For Conditional Noise Contrastive Estimation, we must choose the standard deviation of our noise distribution. We treat this as a hyper-parameter and search over values in [.01, .1, 1.10].

LSD has two hyper-parameters: the number of critic steps for every model step and the \mathcal{L}_2 regularizer λ . We search over 1 and 5 for the critic steps and [.1, 1.10.] for λ . For all dimensions, we find that the using 1 critic step works the best.

Models were all trained for 100,000 steps with the Adam (Kingma & Ba, 2014) optimizer. In both the critic and model, we set the Adam momentum parameters as $\beta_1 = .5, \beta_2 = .9$.

B.6. Image Modeling

Our EBM takes the form

$$q_\theta(x) = \frac{\exp(-E_\theta(x))N(x; \mu, \sigma^2)}{Z} \quad (17)$$

where E_θ is a 2-layer MLP with 1000 hidden units per layer using the Swish nonlinearity. We multiply the energy-function’s output by a Gaussian density (with learned parameters) which guarantees that our model is can be normalized. The critic architecture is identical but with $28 \times 28 = 784$ output dimensions.

We train models on the MNSIT and FashionMNIST datasets for 100 epochs using the Adam optimizer with a step-size of .0001. In both the critic and model, we set the Adam momentum parameters as $\beta_1 = .5, \beta_2 = .9$. We update the critic 5 times for every model update. We set $\lambda = 10$.

B.6.1. SAMPLING

Model samples are drawn with Stochastic Gradient Langevin Dynamics (Welling & Teh, 2011). This is an MCMC sampler with the following update rule:

$$\begin{aligned} x_0 &\sim p(x_0), & \alpha_t &\sim N(0, \epsilon I) \\ x_{t+1} &= x_t - \frac{\epsilon}{2} \nabla_x E_{\theta}(x_t) + \alpha_t \end{aligned} \quad (18)$$

which is quite similar to gradient descent with random Gaussian noise added. To use the sampler in this form, a number of choices must be made. We must select an initialization distribution $p(x_0)$, a step-size ϵ and a number of steps T to run the sampler. Each of these can dramatically affect the distribution of the final sample x_T .

In practice, this sampler can be quite slow to draw samples, so a temperature parameter is added to scale up the impact of the gradient $\nabla_x E_{\theta}(x_t)$ relative to noise α_t . This has the effect of decoupling the step-size from the scale of the Gaussian noise. This changes the update rule to:

$$\begin{aligned} x_0 &\sim p(x_0), & \alpha_t &\sim N(0, \sigma^2 I) \\ x_{t+1} &= x_t - \epsilon \nabla_x E_{\theta}(x_t) + \alpha_t. \end{aligned} \quad (19)$$

This temperature-scaling approach has been used in most recent work on large-scale EBMs (Nijkamp et al., 2019a; Grathwohl et al., 2019; Nijkamp et al., 2019b; Du & Mor-datch, 2019). With this modified version of the sampler we must select: the initialization distribution $p(x_0)$, the step-size ϵ , the noise scale σ and the number of steps T – further complicating the sampling process.

The samples seen in Figure 9 were generated using with $\epsilon = 1$. and $\sigma = .01$ and the sampler was run for 1000 steps. The samples were initialized to a uniform distribution over the data space.

B.6.2. PRE-PROCESSING

The MNIST and Fashion MNIST datasets consist of 28 by 28 images. Each pixel is represented with an integer value in $[0, 1, \dots, 254, 255]$. We dequantize by adding uniform(0, 1) noise to these pixel values, and scale to the range (0, 1). We then apply a logit transformation ($\log x - \log(1 - x)$) which maps $(0, 1) \rightarrow R$. All training and sampling is done in logit-space and samples are mapped back to their original space with a Sigmoid function for visualization.

C. Analysis of the LSD Goodness-of-fit Test Statistic

Our analysis of the test statistic used in the LSD GOF test rely on the central limit theorem taking effect giving our statistic a $N(0, 1)$ distribution under the null hypothesis.

With $n \ll \infty$ one may wonder if our test statistic still has this property. To explore this we run a number of tests to determine normality of data.

To generate statistic samples for testing we run 200 tests under H_0 for RBMs with dimehsion 50, 100, and 200. For each we run a Kolmogorov–Smirnov test and compute a qq-plot. The p -values of the Kolmogorov–Smirnov tests are 0.29, 0.68, and 0.45, respectively indicating that for each model size there is little evidence to indicate the statistic is not distributed as $N(0, 1)$.

Next we show the qq-plot which plots the the quantiles of an empirical histogram against the quantiles of the target distribution ($N(0, 1)$ in this case). Ideally the values should lie on the line $y = x$ indicating the empirical CDF aligns with the theoretical CDF. These plots can be found in Figure 11. As we see, the empirical distribution of the statistic very closely matches the desired distribution for each model size indicating that our test statistic behaves as the theory would suggest.

D. Further Image Modeling Results

D.1. Additional MNIST Samples

We show additional samples from our MNIST model in Figure 12.

D.2. Exploring Sampler Parameters for Deep EBMs

Here we demonstrate the impact of the approximate MCMC sampler’s parameters on sample quality. In Figure 13 we initialize 2 sets of samples to the same value and run SGLD samplers with different noise-scale parameters 300 steps. Clearly, the samples on the right are more diverse and visually appealing, but this does not tell us that they are *better* samples.

Next, we demonstrate the impact of the number of sampler steps on sample quality. We initialize two sets of samples to the same value and run a sampler for 30 steps and then 1000 steps. Results can be seen in Figure 14. We can see that initially the samples are quite diverse and as the sampler is run longer, the samples become more clear and high-quality but considerably less diverse.

We find the choice of sampler parameters has a large impact on the quality and diversity of the resulting samples. Any sample-based evaluation metric would rate each of these sample sets quite differently. Thus these metrics would rate our model differently based on the choice of sampler parameters. Given that these samples are completely separate from our model, this is not ideal behavior.

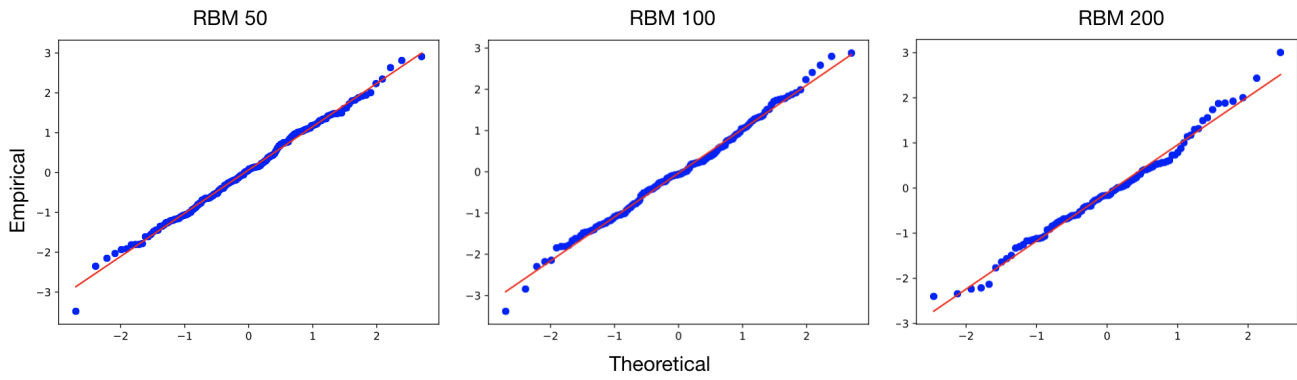


Figure 11. Quantile-Quantile plots of the test statistic compared to $N(0, 1)$ for RBMs of dimension 50, 100, 200, left to right. When the two distributions are identical, the quantiles will fall along the line $y = x$ (red) given enough samples. These results indicated that our test’s p-values are well-calibrated under the null hypothesis.




Figure 12. Additional MNIST samples

D.3. RBMs

We find LSD can be used to train a variety of models – ICA, deep EBMs and also Gaussian-Bernoulli RBMs. Here we train a single-layer RBM with 100 hidden units on MNIST and FashionMNIST. Samples from the models can be seen in Figure 15. Samples were generated with a Gibbs sampler chain run for 2000 iterations.

Models are trained for 100 epochs using the Adam optimizer with a step size of .001. We use 5 critic updates per model update and set $\lambda = 10$. We use the same pre-processing as in our deep EBM experiments.




Figure 13. Samples from a deep EBM with different sampler parameters, but the same initialization. Left: $\sigma = .01$, Right: $\sigma = .1$



Figure 14. Samples from a deep EBM with different sampler parameters but the same initialization. Left: 30 steps, Right: 1000 steps




Figure 15. Samples from RBMs trained with LSD. Left: MNIST, Right: FashionMNIST .