
A Spectral Approach to Gradient Estimation for Implicit Distributions

Jiaxin Shi¹ Shengyang Sun² Jun Zhu¹

Abstract

Recently there have been increasing interests in learning and inference with implicit distributions (i.e., distributions without tractable densities). To this end, we develop a gradient estimator for implicit distributions based on Stein’s identity and a spectral decomposition of kernel operators, where the eigenfunctions are approximated by the Nyström method. Unlike the previous works that only provide estimates at the sample points, our approach directly estimates the gradient function, thus allows for a simple and principled out-of-sample extension. We provide theoretical results on the error bound of the estimator and discuss the bias-variance tradeoff in practice. The effectiveness of our method is demonstrated by applications to gradient-free Hamiltonian Monte Carlo and variational inference with implicit distributions. Finally, we discuss the intuition behind the estimator by drawing connections between the Nyström method and kernel PCA, which indicates that the estimator can automatically adapt to the geometry of the underlying distribution.

1. Introduction

Recently there have been increasing interests in learning and inference with implicit distributions, i.e., distributions defined by a sampling process but without tractable densities. Popular examples include Generative adversarial networks (GAN) (Goodfellow et al., 2014; Mohamed & Lakshminarayanan, 2016). Compared to the explicit likelihoods (e.g., Gaussian) in other deep generative models such as variational autoencoders (VAE) (Kingma & Welling, 2013), implicit distributions are shown able to capture the complex data manifold that lies in a high dimensional space,

leading to more realistic samples generated by GAN than other models. Besides, as the constraint of requiring an explicit density is removed, implicit distributions are treated as more flexible variants of variational families for approximate inference (Ranganath et al., 2016; Liu & Feng, 2016; Mescheder et al., 2017; Tran et al., 2017; Huszár, 2017; Li & Turner, 2018; Shi et al., 2018).

Despite that it is appealing to use flexible implicit distributions, which capture complex correlations and manifold structures, deploying them in practical scenes is still challenging. This is because most learning and inference algorithms require optimizing some divergences between two distributions, which often rely on evaluating the densities of them. However, the density of an implicit distribution is intractable and we only have access to its samples. Previous works have explored two directions to solve the problem. One is to first approximate the optimization objective with the samples and then use the approximation to guide the learning procedure. Many works in this direction are based on the fact that the density ratio between two distributions can be estimated from their samples, by a probabilistic classifier (also known as the discriminator) trained in an adversarial game (Donahue et al., 2016; Dumoulin et al., 2016; Mescheder et al., 2017; Tran et al., 2017; Huszár, 2017), or by kernel-based estimators (Shi et al., 2018). The other direction is to estimate the gradients instead of the objective. Li & Turner (2018) propose the Stein gradient estimator for the log density of an implicit distribution. It is based on a ridge regression that inverts a generalized version of Stein’s identity (Gorham & Mackey, 2015; Liu et al., 2016). As argued in their work, this approach is more direct and avoids probable arbitrarily diverse gradients provided by the approximate objective. In this paper, we focus on the latter direction.

Though the Stein gradient estimator has been shown to be a fast and easy way to obtain gradient estimates for implicit models. It is still limited by its simple formulation, i.e., the unjustified choice of both the test function (the kernel feature mapping) and the regularization scheme (the Frobenius-norm regularization used in ridge regression). The problem has deeper implications. For instance, no theoretical results have been established for the estimator. Moreover, there is no principled way to obtain gradient estimates at positions out of the sample points.

¹Dept. of Comp. Sci. & Tech., BNRist Center, State Key Lab for Intell. Tech. & Sys., THBI Lab, Tsinghua University
²Dept. of Comp. Sci., University of Toronto. Correspondence to: Jiaxin Shi <shijx15@mails.tsinghua.edu.cn>, Jun Zhu <dc-szj@tsinghua.edu.cn>.

In this paper, we develop a novel gradient estimator for implicit distributions, which is called the *Spectral Stein Gradient Estimator* (SSGE). To approximate the gradient function of the log density (i.e., $\nabla_{\mathbf{x}} \log q(\mathbf{x})$), SSGE expands it in terms of the eigenfunctions of a kernel-based operator. These eigenfunctions are orthogonal with respect to the underlying distribution. By setting the test functions in the Stein’s identity to these eigenfunctions, we can take advantage of their orthogonality to obtain a simple solution. The eigenfunctions in the solution are then approximated by the Nyström method (Nyström, 1930; Baker, 1997; Williams & Seeger, 2001). Unlike the Stein gradient estimator (Li & Turner, 2018), our approach allows for a direct and principled out-of-sample extension. Moreover, we provide theoretical analysis on the error bound of SSGE and discuss the bias-variance tradeoff in practice. We also discuss its effectiveness in reducing the curse of dimensionality by drawing connections to kernel PCA.

2. Background

In this section we briefly introduce the Nyström method and the Stein gradient estimator.

2.1. The Nyström Method

The Nyström method originates as a method for approximating the solution of Fredholm integral equations of the second kind (Nyström, 1930; Baker, 1997). It was used by Williams & Seeger (2001) for estimating extensions of eigenvectors in Gaussian process regression. Specifically, the following equation for finding the eigenfunctions $\{\psi_j\}_{j \geq 1}$, $\psi_j \in L^2(\mathcal{X}, q)^1$ of the covariance kernel $k(\mathbf{x}, \mathbf{y})$ w.r.t. the probability measure q is considered:

$$\int k(\mathbf{x}, \mathbf{y}) \psi(\mathbf{y}) q(\mathbf{y}) d\mathbf{y} = \mu \psi(\mathbf{x}). \quad (1)$$

And there is a constraint that the eigenfunctions $\{\psi_j\}_{j \geq 1}$ are orthonormal under q :

$$\int \psi_i(\mathbf{x}) \psi_j(\mathbf{x}) q(\mathbf{x}) d\mathbf{x} = \delta_{ij}, \quad (2)$$

where $\delta_{ij} = \mathbb{1}[i = j]$. Approximating the left side of eq. (1) with its unbiased Monte Carlo estimate using i.i.d. samples $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ from q and applying the equation to these samples, we obtain

$$\frac{1}{M} \mathbf{K} \boldsymbol{\psi} \approx \mu \boldsymbol{\psi}, \quad (3)$$

where \mathbf{K} is the Gram matrix: $\mathbf{K}_{ij} = k(\mathbf{x}^i, \mathbf{x}^j)$, and $\boldsymbol{\psi} = [\psi(\mathbf{x}^1), \dots, \psi(\mathbf{x}^M)]^\top$. This is an eigenvalue problem for \mathbf{K} . We compute the eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_J$ with the J

¹ $L^2(\mathcal{X}, q)$ denotes the space of all square-integrable functions w.r.t. q .

largest eigenvalues $\lambda_1 \geq \dots \geq \lambda_J$ for \mathbf{K} . Now we have the solutions of eq. (3) by comparing against $\mathbf{K} \mathbf{u}_j = \lambda_j \mathbf{u}_j$:

$$\psi_j(\mathbf{x}^m) \approx \sqrt{M} u_{jm}, \quad m = 1, \dots, M, \quad (4)$$

$$\mu_j \approx \frac{\lambda_j}{M}. \quad (5)$$

Note that the scaling factor in eq. (4) is due to the empirical constraint translated from eq. (2): $\frac{1}{M} \sum_{m=1}^M \psi_i(\mathbf{x}^m) \psi_j(\mathbf{x}^m) \approx \delta_{ij}$. Baker (1997) shows that for a fixed kernel k , $\frac{\lambda_j}{M}$ converges to μ_j in the limit as $M \rightarrow \infty$.

Plugging these solutions back into eq. (1), we get the Nyström formula for approximating the value of the j th eigenfunction at any point \mathbf{x} :

$$\psi_j(\mathbf{x}) \approx \hat{\psi}_j(\mathbf{x}) = \frac{\sqrt{M}}{\lambda_j} \sum_{m=1}^M u_{jm} k(\mathbf{x}, \mathbf{x}^m). \quad (6)$$

The Nyström method has been shown to be a thread linking many dimension reduction methods such as kernel PCA (Schölkopf et al., 1998), multidimensional scaling (MDS) (Borg & Groenen, 2005), local linear embedding (LLE) (Roweis & Saul, 2000), Laplacian eigenmaps (Belkin & Niyogi, 2003), and spectral clustering (Weiss, 1999), unifying their out-of-sample extensions (Bengio et al., 2004a;b; Burges et al., 2010). We will discuss these connections later.

2.2. Stein’s Identity and Stein Gradient Estimator

Recent developments in Stein discrepancy and its kernelized extensions (Gorham & Mackey, 2015; Chwialkowski et al., 2016; Liu et al., 2016; Liu & Wang, 2016) have renewed the interests in Stein’s method, which is a classic tool in statistics. Central to these works is an equation that generalizes the original Stein’s identity (Stein, 1981), shown in the following theorem.

Theorem 1 (Gorham & Mackey 2015; Liu et al. 2016). *Assume that $q(\mathbf{x})$ is a continuous differentiable probability density supported on $\mathcal{X} \subset \mathbb{R}^d$. $\mathbf{h} : \mathcal{X} \rightarrow \mathbb{R}^{d'}$ is a smooth vector-valued function $\mathbf{h}(\mathbf{x}) = [h_1(\mathbf{x}), h_2(\mathbf{x}), \dots, h_{d'}(\mathbf{x})]^\top$, and $\forall i \in 1, \dots, d'$, h_i is in the **Stein class** of q , i.e.,*

$$\int_{\mathbf{x} \in \mathcal{X}} \nabla_{\mathbf{x}} (h_i(\mathbf{x}) q(\mathbf{x})) d\mathbf{x} = 0. \quad (7)$$

Then the following identity holds:

$$\mathbb{E}_q[\mathbf{h}(\mathbf{x}) \nabla_{\mathbf{x}} \log q(\mathbf{x})^\top + \nabla_{\mathbf{x}} \mathbf{h}(\mathbf{x})] = \mathbf{0}. \quad (8)$$

The condition (7) can be easily checked using integration by parts or divergence theorem. Specifically, when $\mathcal{X} = \mathbb{R}^d$, eq. (7) holds if $\lim_{\|\mathbf{x}\| \rightarrow \infty} \mathbf{h}(\mathbf{x}) q(\mathbf{x}) = \mathbf{0}$; when \mathcal{X} is a compact subset of \mathbb{R}^d with piecewise smooth boundary $\partial \mathcal{X}$, eq. (7) holds if $\mathbf{h}(\mathbf{x}) q(\mathbf{x}) = \mathbf{0}, \forall \mathbf{x} \in \partial \mathcal{X}$. Here $\mathbf{h}(\mathbf{x})$ is

called the test function. We can check that for a RBF kernel k , and for any fixed \mathbf{x} , $k(\mathbf{x}, \cdot)$ and $k(\cdot, \mathbf{x})$ are in the Stein class of continuous differentiable densities supported on \mathbb{R}^d .

Because the expectation in eq. (8) can be approximated by Monte Carlo estimates, the identity connects $\nabla_{\mathbf{x}} \log q(\mathbf{x})$ and the samples from q . Inspired by this, Li & Turner (2018) propose the Stein gradient estimator, which inverts eq. (8) to obtain estimates of $\nabla_{\mathbf{x}} \log q(\mathbf{x})$ at the sample points. Below we briefly review their method. Specifically, consider M i.i.d. samples $\mathbf{x}^{1:M}$ drawn from $q(\mathbf{x})$. We define two matrices $\mathbf{H} = [\mathbf{h}(\mathbf{x}^1), \dots, \mathbf{h}(\mathbf{x}^M)] \in \mathbb{R}^{d' \times M}$ and $\mathbf{G} = [\nabla_{\mathbf{x}^1} \log q(\mathbf{x}^1), \dots, \nabla_{\mathbf{x}^M} \log q(\mathbf{x}^M)]^\top \in \mathbb{R}^{M \times d}$. Monte Carlo sampling with eq. (8) shows that

$$-\frac{1}{M} \mathbf{H} \mathbf{G} \approx \overline{\nabla_{\mathbf{x}} \mathbf{h}}, \quad (9)$$

where $\overline{\nabla_{\mathbf{x}} \mathbf{h}} = \frac{1}{M} \sum_{m=1}^M \nabla_{\mathbf{x}^m} \mathbf{h}(\mathbf{x}^m) \in \mathbb{R}^{d' \times d}$, $\nabla_{\mathbf{x}^m} \mathbf{h}(\mathbf{x}^m) = [\nabla_{\mathbf{x}^m} h_1(\mathbf{x}^m), \dots, \nabla_{\mathbf{x}^m} h_{d'}(\mathbf{x}^m)]^\top$. Equation (9) inspires the following ridge regression problem:

$$\operatorname{argmin}_{\hat{\mathbf{G}} \in \mathbb{R}^{M \times d}} \|\overline{\nabla_{\mathbf{x}} \mathbf{h}} + \frac{1}{M} \mathbf{H} \hat{\mathbf{G}}\|_F^2 + \frac{\eta}{M^2} \|\hat{\mathbf{G}}\|_F^2,$$

where $\|\cdot\|_F^2$ denotes the Frobenius norm of a matrix and $\eta > 0$ is the regularization coefficient. It has an analytic solution that

$$\hat{\mathbf{G}}^{\text{Stein}} = -M(\mathbf{K} + \eta \mathbf{I})^{-1} \mathbf{H}^\top \overline{\nabla_{\mathbf{x}} \mathbf{h}}, \quad (10)$$

where $\mathbf{K} = \mathbf{H}^\top \mathbf{H}$. By noticing that $\mathbf{K}_{ij} = \mathbf{h}(\mathbf{x}^i)^\top \mathbf{h}(\mathbf{x}^j)$ and applying the kernel trick, we have $\mathbf{K}_{ij} = k(\mathbf{x}^i, \mathbf{x}^j)$, where $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a positive definite kernel. Similarly we can show $(\mathbf{H}^\top \overline{\nabla_{\mathbf{x}} \mathbf{h}})_{ij} = \frac{1}{M} \sum_{m=1}^M \nabla_{\mathbf{x}^m} k(\mathbf{x}^i, \mathbf{x}^m)$. With the kernel trick, the above derivation implicitly set the test function to be the feature mapping $\mathbf{h} : \mathbb{R}^d \rightarrow \mathcal{H}$, $\mathbf{h}(\mathbf{x}) = k(\mathbf{x}, \cdot)$, where \mathcal{H} is the Reproducing Kernel Hilbert Space (RKHS) induced by k .

Though introducing the kernel trick enhances the expressiveness of the Stein gradient estimator, the choice of the test function is not well justified. It is unclear whether eq. (8) holds when the test function is set to the kernel feature mapping $\mathbf{h}(\mathbf{x}) = k(\mathbf{x}, \cdot)$, which maps from \mathbb{R}^d to \mathcal{H} instead of the common $\mathbb{R}^{d'}$. Li & Turner (2018) show that their approach is equivalent to minimizing a regularized version of the V-statistics of kernel Stein discrepancy (Liu et al., 2016), which has proved to be effective in testing goodness-of-fit. However, it is still unclear whether the test power is sufficient due to the added Frobenius-norm regularization term. Besides, eq. (10) only gives the gradient estimates at the sample points. For out-of-sample prediction at a test point, two choices are proposed in Li & Turner (2018): one is by adding the test point to the sample points and re-compute eq. (10), which could be computationally

demanding. We will refer to this out-of-sample extension as *Stein⁺*. The other choice is to refit a parametric estimator (a linear combination of RBF kernels), which is cheaper but less accurate. Both approaches assume that the test points are also sampled from q . However, they are unjustified when the assumption is not satisfied. Since the latter is an approximation to the former, we will only compare with *Stein⁺* in the experiments for out-of-sample predictions.

3. Method

In this section we derive a gradient estimator for implicit distributions called the *Spectral Stein Gradient Estimator* (SSGE). Unlike the previous Stein gradient estimator that only provides estimates at the sample points, SSGE directly estimates the gradient function and thus allows simple and principled out-of-sample predictions. We also provide theoretical analysis on the error bound of SSGE.

3.1. Spectral Stein Gradient Estimator

To begin with, let \mathbf{x} denote a d -dimensional vector in \mathbb{R}^d . Consider an implicit distribution $q(\mathbf{x})$ supported on $\mathcal{X} \subset \mathbb{R}^d$, from which we observe M i.i.d. samples $\mathbf{x}^{1:M}$. We denote the target gradient function to estimate by $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^d$: $\mathbf{g}(\mathbf{x}) = \nabla_{\mathbf{x}} \log q(\mathbf{x})$. The i th component of the gradient is $g_i(\mathbf{x}) = \nabla_{x_i} \log q(\mathbf{x})$. We assume $g_1, \dots, g_d \in L^2(\mathcal{X}, q)$. As introduced in section 2.1, $\{\psi_j\}_{j \geq 1}$ form an orthonormal basis of $L^2(\mathcal{X}, q)$. So we can expand $g_i(\mathbf{x})$ into the following spectral series:

$$g_i(\mathbf{x}) = \sum_{j=1}^{\infty} \beta_{ij} \psi_j(\mathbf{x}). \quad (11)$$

Below we will show how to estimate the coefficients β_{ij} . According to Theorem 1, we have the following proposition.

Proposition 1. *If $k(\cdot, \cdot)$ has continuous second order partial derivatives, and both $k(\mathbf{x}, \cdot)$ and $k(\cdot, \mathbf{x})$ are in the Stein class of q , the following set of equations hold true:*

$$\mathbb{E}_q[\psi_j(\mathbf{x})g(\mathbf{x}) + \nabla_{\mathbf{x}} \psi_j(\mathbf{x})] = \mathbf{0}, \quad j = 1, 2, \dots, \infty. \quad (12)$$

Proof. We only need to prove that $\psi_j(\mathbf{x})$ is in the Stein class of q . See Appendix A for details. \square

Substituting eq. (11) into eq. (12) and using the orthonormality of $\{\psi_j\}_{j \geq 1}$, we can show that

$$\beta_{ij} = -\mathbb{E}_q \nabla_{x_i} \psi_j(\mathbf{x}).$$

To estimate β_{ij} , we need an approximation of $\nabla_{x_i} \psi_j(\mathbf{x})$. The key observation is that derivatives can be taken w.r.t. both sides of eq. (1):

$$\begin{aligned} \mu_j \nabla_{x_i} \psi_j(\mathbf{x}) &= \nabla_{x_i} \int k(\mathbf{x}, \mathbf{y}) \psi_j(\mathbf{y}) q(\mathbf{y}) d\mathbf{y} \\ &= \int \nabla_{x_i} k(\mathbf{x}, \mathbf{y}) \psi_j(\mathbf{y}) q(\mathbf{y}) d\mathbf{y}. \end{aligned} \quad (13)$$

Monte-Carlo sampling with eq. (13), we have an estimate of $\nabla_{x_i} \psi_j(\mathbf{x})$:

$$\hat{\nabla}_{x_i} \psi_j(\mathbf{x}) \approx \frac{1}{\mu_j M} \sum_{m=1}^M \nabla_{x_i} k(\mathbf{x}, \mathbf{x}^m) \psi_j(\mathbf{x}^m). \quad (14)$$

Substituting eqs. (4) and (5) into eq. (14) and comparing with eq. (6), we can show

$$\hat{\nabla}_{x_i} \psi_j(\mathbf{x}) \approx \nabla_{x_i} \hat{\psi}_j(\mathbf{x}). \quad (15)$$

Perhaps surprisingly, Equation (15) indicates that $\nabla_{x_i} \hat{\psi}_j(\mathbf{x})$ is a good approximation to $\nabla_{x_i} \psi_j(\mathbf{x})^2$. In fact, as we shall see in Theorem 2, the error introduced by Nyström approximation is negligible with high probability as $M \rightarrow \infty$.

Now truncating the series expansion to the first J terms and plugging in the Nyström approximations of $\{\psi_j\}_{j=1}^J$, we get our estimator:

$$\hat{g}_i(\mathbf{x}) = \sum_{j=1}^J \hat{\beta}_{ij} \hat{\psi}_j(\mathbf{x}), \quad (16)$$

$$\hat{\beta}_{ij} = -\frac{1}{M} \sum_{m=1}^M \nabla_{x_i} \hat{\psi}_j(\mathbf{x}^m), \quad (17)$$

where $\hat{\psi}_j$ is the Nyström approximation of ψ_j as in Section 2.1. We use RBF kernels in all experiments.

Computational Cost The spectral gradient estimator $\hat{g}_i(\mathbf{x})$ depends on $\hat{\beta}_{ij}$ and $\hat{\psi}_j(\mathbf{x})$. To compute them, the computational bottleneck lies in computing the Gram matrix and its eigendecomposition, which have complexity $O(M^2 d)$ and $O(M^3)$, respectively. Therefore the computational cost of constructing the estimator is $O(M^3 + M^2 d)$. For prediction, given $\mathbf{x} \in \mathbb{R}^d$, evaluating $\hat{g}_i(\mathbf{x})$ has cost $O(M(d + J))$. In comparison, the Stein gradient estimator directly approximates gradients at the sample points, which involves computing the Gram matrix and its inverse. Thus the overall complexity is also $O(M^3 + M^2 d)$. Note that SSGE only requires the J largest eigenvalues and corresponding eigenvectors, efficient algorithms (Parlett & Scott, 1979) might be applied to further reduce its complexity.

3.2. Theoretical Results

Following the derivation in Section 3.1, we analyze the theoretical properties of the resulting estimator in eqs. (16) and (17). To be clear, we formally restate the assumptions that have been made in the derivation.

Assumption 1. $k(\mathbf{x}, \cdot)$ and $k(\cdot, \mathbf{x})$ are in the Stein class of q .

Assumption 2. $g_i(\mathbf{x}) \in L^2(\mathcal{X}, q)$, $i = 1, \dots, d$, i.e.,

$$\int g_i(\mathbf{x})^2 q(\mathbf{x}) d\mathbf{x} = \sum_{j=1}^{\infty} \beta_{ij}^2 \leq C < \infty.$$

²This does not hold for general functions.

Assumption 3. $\mu_1 > \mu_2 > \dots > \mu_J > 0$.

Note that Assumption 1 holds for RBF kernels. Assumption 2 is necessary for $g_i(\mathbf{x})$'s being possible to be expanded into the spectral series. We need Assumption 3 since our derivation is based on several well-studied bounds of Nyström approximation, i.e., Lemmas 4 and 5 in Appendix B (Sinha & Belkin, 2009; Izbicki et al., 2014). Note that when this assumption does not hold, we could proceed as in Rosasco et al. (2010) (Theorem 12) and derive the error bound in a similar way.

Theorem 2 (Error Bound of SSGE, proof in Appendix B). *Given the above assumptions, the error $\int |\hat{g}_i(\mathbf{x}) - g_i(\mathbf{x})|^2 q(\mathbf{x}) d\mathbf{x}$ is bounded by*

$$J^2 \left(O_p \left(\frac{1}{M} \right) + O_p \left(\frac{C}{\mu_J \Delta_J^2 M} \right) \right) + JO_p \left(\frac{C}{\mu_J \Delta_J^2 M} \right) + \|g_i\|_{\mathcal{H}}^2 O(\mu_J), \quad (18)$$

where $\Delta_J = \min_{1 \leq j \leq J} |\mu_j - \mu_{j+1}|$, O_p is the Big O notation in probability.

The first three terms in eq. (18) are the sample errors caused by the Nyström approximation, which we call the *estimation error*. It is negligible with high probability as $M \rightarrow \infty$. The last term is caused by the bias introduced by the truncation, which we call the *approximation error*. From the bound we can observe a tradeoff between the estimation error and the approximation error. As an illustration, one may set J to be as large as possible to reduce the magnitude of μ_J and thus reduce the approximation error, but it will increase the estimation error at a rate of $O_p \left(\frac{J^2}{\mu_J} \right)$.

For RBF kernels and their corresponding RKHS, smoother target functions tend to have smaller $\|g_i\|_{\mathcal{H}}^2$, and thus a tighter bound. In general, this indicates that choosing the appropriate kernel which is suitable to the target gradient function can improve the performance of the gradient estimator (by leading to a smaller $\|g_i\|_{\mathcal{H}}^2$).

Hyperparameter Selection When RBF kernels are used, SSGE has two free parameters: The kernel bandwidth σ , and the number of eigenfunctions (J) used in the estimate. For σ , we use the median heuristic, i.e., we set it to be the median of pairwise distances between all samples, which turns out to work well in all experiments. Below we discuss the criterion for selecting J , which is usually harder.

As the performance of the gradient estimator directly influences the task where it is used. The optimal choice for tuning J is to apply cross-validation on the specific task. However, since J is a discrete parameter, one has to manually set a continuous interval and bin it so that the commonly used black-box hyperparameter-search methods (e.g., Bayesian

optimization) are applicable. However, this approach does not take the magnitude of eigenvalues into consideration. Observing this, we propose that, instead of directly tuning J , we could tune a threshold \bar{r} for the percentage of remaining eigenvalues:

$$J = \operatorname{argmax}_{J'} r_{J'}, \text{ s.t. } r_{J'} = \frac{\sum_{j=1}^{J'} \lambda_j}{\sum_{j=1}^M \lambda_j}, r_{J'} \leq \bar{r}.$$

Note that searching \bar{r} may still not be easy due to the non-smooth validation surface. But in experiments we found that \bar{r} values in $[0.95, 0.99]$ usually work well.

3.3. Gradient Estimation for Entropy

Above we have derived a gradient estimator for the log density of an implicit distribution. Now we discuss a useful extension of it. Consider the situation where we need to optimize the entropy $\mathbb{H}(q) = -\mathbb{E}_q \log q$ of an implicit distribution $q_\phi(\mathbf{x})$ w.r.t. its parameters ϕ . When \mathbf{x} is continuous and can be reparameterized (Kingma & Welling, 2013), e.g., $\mathbf{x} = f(\epsilon; \phi)$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, it can be shown (see Appendix C) that

$$\nabla_\phi \mathbb{H}(q) = -\mathbb{E}_\epsilon \nabla_{\mathbf{x}} \log q_\phi(f(\epsilon; \phi)) \nabla_\phi f(\epsilon; \phi), \quad (19)$$

where $\nabla_{\mathbf{x}} \log q_\phi(f(\epsilon; \phi))$ can be easily estimated by SSGE. As we shall see in experiments, eq. (19) can be used for variational inference with implicit distributions.

4. Related Work

Our work is closely related to other works on implicit distributions. Apart from the density-ratio based approaches introduced in Section 1, we discuss two more directions.

Nonparametric Inference Nonparametric variational inference (VI) methods such as PMD (Dai et al., 2016) and SVGD (Liu & Wang, 2016) remove the need of parametric families, they keep a set of particles and gradually adjust them towards the true posterior. These particles can be viewed as samples from an implicit distribution. Instead of directly computing gradients in the sample space like us, SVGD performs functional gradient descent to transform the implicit distribution towards the true posterior. Though elegant, SVGD is limited to KL-divergence based VI problems, while our approach is generally applicable wherever gradient estimates are needed for the log density of implicit distributions.

Kernel Exponential Families and Score Matching Previous to our work, the problem of estimating gradient functions of intractable log densities has been worked on by Strathmann et al. (2015); Sutherland et al. (2018). They identified the problem when developing Hamiltonian Monte Carlo (HMC) under the settings where higher-order information of the target distribution is unavailable. To address




Figure 1. Gradient estimates of the log density of $\mathcal{N}(0, 1)$.

it, a kernel exponential family (Sriperumbudur et al., 2017) is fit to samples along the Markov Chain trajectory by score matching (Hyvärinen, 2005), and then serves as a surrogate of the target distribution to provide gradient estimates for HMC. We will compare to them in Section 5.2.

5. Experiments

We evaluate the proposed approach on both toy problems and real-world examples. The latter includes applications of SSGE to two widely used inference methods: Hamiltonian Monte Carlo and variational inference. Code is available at <https://github.com/thjashin/spectral-stein-grad>. Implementations are based on ZhuSuan (Shi et al., 2017).

5.1. Toy Experiment

As a simple example, we experiment with estimating the gradient function of a 1-D standard Gaussian distribution. The target log density is $\log q(x) = -\frac{1}{2} \log 2\pi - \frac{1}{2}x^2$, and the true gradient function is $\nabla_x \log q(x) = -x$. We draw $M = 100$ i.i.d. samples from q for use in the estimation. In Figure 1 we plot the gradients estimates produced by the Stein gradient estimator, its out-of-sample extension (Stein⁺) (see Section 2.2), and our approach (SSGE). Since the original Stein estimator only gives estimates at the sample points, we plot them as individual points (in red). For the regularization coefficient η in eq. (10), we searched it in $\{0.001, 0.01, 0.1, 1, 10, 100\}$ and plot the best result at $\eta = 0.1$ ³. For SSGE, we set $J = 6$. We can see that despite all three estimators produce rather good approximation where the samples are taken densely (e.g., in $[-2, 2]$), the gradient function estimated by SSGE is notably better at the places where samples are less dense.

5.2. Gradient-free Hamiltonian Monte Carlo

In this experiment we investigate the usefulness of SSGE in constructing a gradient-free HMC sampler. We follow the

³The criterion for selecting η is unclear in Li & Turner (2018).

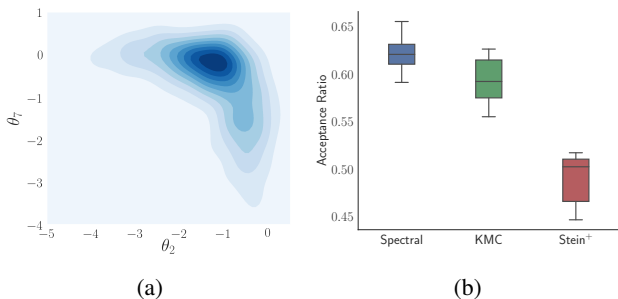


Figure 2. (a) Dimensions 2 and 7 of the marginal hyperparameter posterior of Gaussian Process classification on the UCI Glass dataset; (b) The average acceptance ratios of gradient-free HMC using SSGE, KMC, and Stein⁺.

settings in Sejdinovic et al. (2014); Strathmann et al. (2015) and consider a Gaussian Process classification problem on the UCI Glass dataset. The goal is to infer the posterior over hyperparameters under a fully Bayesian treatment. Specifically, consider a Gaussian process whose joint distribution is over latent variables \mathbf{f} , labels \mathbf{y} , and hyperparameters θ :

$$p(\mathbf{f}, \mathbf{y}, \theta) = p(\theta)p(\mathbf{f}|\theta)p(\mathbf{y}|\mathbf{f}), \quad (20)$$

where $\mathbf{f}|\theta \sim N(\mathbf{0}, \mathbf{K}_\theta)$. \mathbf{K}_θ is the Gram matrix formed by the data points $\mathbf{x}_{1:N} \in \mathbb{R}^D$:

$$(\mathbf{K}_\theta)_{ij} = \exp \left\{ - \sum_{d=1}^D \frac{|x_{i,d} - x_{j,d}|^2}{2\ell_d^2} \right\}, \quad (21)$$

where we define $\theta_d = \log \ell_d^2$. The problem to consider is a binary classification between window and non-window glasses, so the likelihood is given by a logistic classifier: $p(y_i|f_i) = \frac{1}{1 + \exp(-y_i f_i)}$, $y_i \in \{-1, 1\}$. The posterior over θ is highly nonlinear, as shown in Figure 2a. As pointed out in previous works (Murray & Adams, 2010; Sejdinovic et al., 2014), sampling from the posterior of θ is challenging, e.g., Gibbs sampling often gets stuck due to $p(\theta|\mathbf{f}, \mathbf{y})$ is very sharp. One way to address this problem is the pseudo-marginal MCMC (Andrieu et al., 2009), through which a markov chain can be simulated to directly sample from $p(\theta|\mathbf{y})$. Since the likelihood $p(\mathbf{y}|\theta)$ is intractable, pseudo-marginal MCMC replaces it with an unbiased Monte Carlo estimate using importance sampling:

$$\hat{p}(\mathbf{y}|\theta) = \frac{1}{K} \sum_{i=1}^K \frac{p(\mathbf{y}|\mathbf{f}^i)p(\mathbf{f}^i|\theta)}{q(\mathbf{f}^i)}, \quad \mathbf{f}^{1:K} \sim q(\mathbf{f}). \quad (22)$$

In practice $q(\mathbf{f})$ is chosen to be the Laplace approximation of $p(\mathbf{f}|\mathbf{y}, \theta)$. As for any pseudo-marginal MCMC scheme, the gradient information of the posterior is not available and HMC is not suitable. We have to resort to gradient-free MCMC methods. As mentioned in Section 4, kernel adaptive Metropolis samplers (Sejdinovic et al., 2014) were

developed and then extended to kernel HMC (KMC) (Strathmann et al., 2015). In this experiment we compare the performance of KMC and HMC with gradients estimated by SSGE and Stein⁺.

To begin, we run 20 randomly initialized adaptive-Metropolis samplers for 30k iterations, with the first 10k samples discarded. We then keep every 400-th sample in each of the chains, and combine them to get 1k samples. These samples are treated as the ground truth. Similar to Sutherland et al. (2018), our experiment assumes the idealized scenario where a burn-in period for collecting a sufficient number of samples has completed. This is to remove all the other factors that could have an effect on the comparison of acceptance ratios, which then only depends on the accuracy of the gradient estimation of potentials, i.e., $-\nabla_\theta \log p(\theta|\mathbf{y})$. So we fit all three estimators on a random subset of $M = 200$ of these samples and repeated 10 times. For each fitted estimator, we start from a random initial point from the posterior sketch and run HMC samplers with the gradient estimates for 5k iterations. To be fair in comparing the acceptance ratios, no adaptation of HMC parameters can be used. So we randomly uses between 1 and 10 leapfrog steps of size chosen uniformly in $[0.01, 0.1]$, and a standard Gaussian momentum. The kernel bandwidths (σ) in all three estimators are determined by median heuristics (i.e., set to the median of the pairwise distances between the M data points). For Stein⁺, $\eta = 0.001$. For SSGE, $\bar{r} = 0.95$.

The average acceptance ratios over 10 runs are plotted in Figure 2b. We can see that SSGE clearly outperforms Stein⁺ and is even better than the KMC algorithm, which is specially designed as a gradient-free HMC algorithm. Though KMC is carefully designed, its gradients are estimated by first fitting a kernel exponential family as a surrogate and then taking derivatives through it, while SSGE is arguably more direct.

5.3. Variational Inference with Implicit Distributions

As introduced in Section 1, there have been increasing interests in constructing flexible variational posteriors with implicit distributions. Specifically, for a latent-variable model $p(\mathbf{z}, \mathbf{x})$ where \mathbf{x} and \mathbf{z} denote observed and latent variables, respectively, Variational Inference (VI) approximates the posterior $p(\mathbf{z}|\mathbf{x})$ by maximizing the following evidence lower bound (ELBO):

$$\mathcal{L}(\mathbf{x}; \phi) = \mathbb{E}_{q_\phi(\mathbf{z})} \log p(\mathbf{z}, \mathbf{x}) - \mathbb{E}_{q_\phi(\mathbf{z})} \log q_\phi(\mathbf{z}), \quad (23)$$

where $q_\phi(\mathbf{z})$ is called the variational distribution. The second term of eq. (23) is the entropy of q , which is intractable for implicit distributions. As shown in Section 3.3, SSGE can be used here for estimating gradients of the entropy term, thus allowing VI with implicit distributions. Below we conduct experiments on two examples: Bayesian Neural Networks (BNN) and Variational Autoencoders (VAE).




Figure 3. Prediction results for the 1-D regression experiment. Spectral-Implicit and Spectral-Factorized represent using SSGE to perform VI with implicit posteriors and factorized Gaussian posteriors, respectively. Shaded areas represent 3 times standard deviation.

Note that the original Stein gradient estimator can also be used here. In experiments, we find that despite lack of theoretical evidences, the performance of a well-tuned Stein gradient estimator is very close to SSGE when no out-of-sample predictions are required. As the emphasis of this paper is on SSGE, we only focus on verifying the accuracy of SSGE below.

5.3.1. BAYESIAN NEURAL NETWORKS

We evaluate the predictive ability of BNNs with implicit variational posteriors trained by SSGE. To visually access the quality of uncertainty, we choose a 1-D regression problem (Hernández-Lobato & Adams, 2015; Louizos & Welling, 2016). Specifically, 20 inputs are randomly sampled from $[-4, 4]$, then the target value y is computed with $y = x^3 + \epsilon_n$, $\epsilon_n \sim \mathcal{N}(0, 9)$. We use a BNN with 1 hidden layer and 20 units to model the normalized inputs and targets. We also set the variance of the observation noise to the true value. We compare SSGE with implicit posteriors, Hamiltonian Monte Carlo (HMC) (Neal et al., 2011) and Bayes-by-backprop (BBB) (Blundell et al., 2015). To better demonstrate SSGE’s gradient estimation effect, we also test SSGE with a factorized posterior, in comparison to BBB.

We keep 20 chains and run 100k iterations for HMC. All other methods are trained with 100 samples for 20k iterations using Adam optimizer (Kingma & Ba, 2014). For SSGE, we set $J = 100$. The implicit posteriors we use for weights in both layers are standard normal distributions transformed by fully connected networks with one hidden layer of 100 units.

As shown in Fig. 3, HMC, as the golden standard, smoothly fits the training data and outputs sensible uncertainty esti-

mation. HMC not only produces large uncertainty outside the data region, its predictive variance also varies even in regions with training points. This kind of interpolation behavior is hard to be captured by factorized Gaussian posteriors, as shown in Fig. 3c and 3d. SSGE with implicit posteriors also has big predictive variances beyond training points, which implies that the BNN trained by SSGE is not overfitting, although it underestimates the uncertainty in the middle region. Also, we observe that SSGE can have similar interpolation behaviors as HMC (see the rightmost two points in Fig. 3b). Besides, SSGE with a factorized posterior has a similar prediction with BBB. Given the network and the variational posterior are both the same, we can attribute this similarity to accurate gradient estimation by SSGE.

5.3.2. VARIATIONAL AUTOENCODERS

From the above example we see that SSGE enables variational posteriors parameterized by implicit distributions. To demonstrate that it scales to larger models and datasets, we adopt the settings in Shi et al. (2018) and train a deep convolutional VAE with implicit variational posteriors (*Implicit VAE* for short) on the CelebA dataset. As in their work, the latent dimension is 32, and the network structure of the decoder is chosen to be the same as DCGAN (Radford et al., 2015). The observation likelihoods are Gaussian distributions with trainable data-independent variances. The implicit posterior is a deep convolutional net symmetric to the decoder, with Gaussian noises injected into hidden layers. Full details of the model structures can be found in Shi et al. (2018).

To examine how accurate the gradient estimates provided by SSGE are, we conduct experiments under three different settings: a plain VAE with normal variational posteriors, an Implicit VAE trained with the entropy term removed from the ELBO, and an Implicit VAE using SSGE’s gradient estimates for the entropy. For SSGE, we set $M = 100$, and $\bar{r} = 0.99$. In Figures 4a to 4c we show samples randomly generated from the trained models. We can see that without the entropy term, the Implicit VAE tends to overfit and produces visually bad generations, while if we retain the entropy term and use SSGE to estimate its gradients, the Implicit VAE can produce realistic samples. To quantitatively measure the sample quality, we compare the Fréchet Inception Distance (FID) (Heusel et al., 2017) between real data and random generations from the models. The results are shown in Figure 4d. We can see that the Implicit VAE trained by SSGE converges faster and produces samples with slightly better quality than the plain VAE. This is probably due to that implicit posteriors are less likely to overfit (Shi et al., 2018), and SSGE gives accurate gradients for optimizing them. Besides the CelebA experiments, we also tested the models on MNIST dataset and evaluated the test log likelihoods. See Appendix D for details.

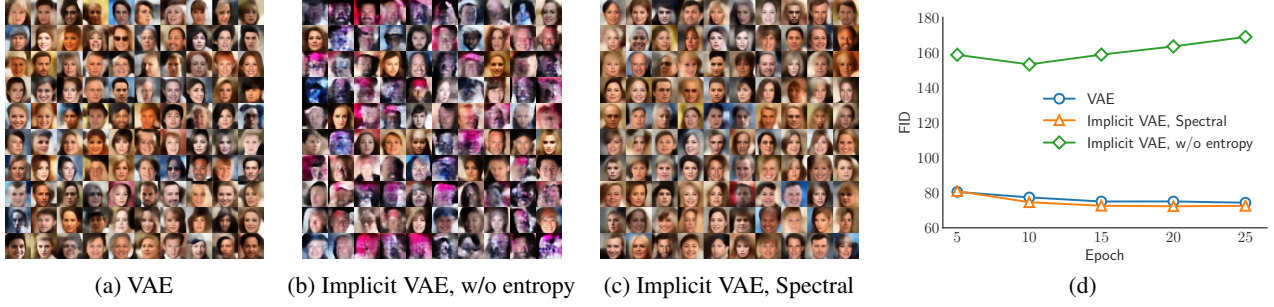


Figure 4. (a)-(c) CelebA samples generated by VAE, Implicit VAE trained without the entropy term, and Implicit VAE trained by SSGE; (d) Fréchet Inception Distances (FID) (Heusel et al., 2017) between random generated images and real images.

6. Discussion

Connection to Kernel PCA As mentioned in Section 2.1, the Nyström approximation is closely related to Kernel PCA (Schölkopf et al., 1998) (KPCA), which is a powerful method for nonlinear dimension reduction. In KPCA, the input data is first projected to a (usually high-dimensional) feature space, where PCA is then applied. The operations in the feature space are handled by the kernel trick. We briefly review the method below.

Given a positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, we denote the induced RKHS by \mathcal{H} and its corresponding feature map by $\phi : \mathcal{X} \rightarrow \mathcal{H}$. Let the data to be analyzed be $\{\mathbf{x}^i\}_{i=1}^M$, $\mathbf{x}^i \in \mathcal{X}$. To simplify the derivation, we first assume the data to be centered in the feature space. Then the covariance matrix is formed as $\mathbf{C} = \frac{1}{M} \sum_{i=1}^M \phi(\mathbf{x}^i) \phi(\mathbf{x}^i)^\top$. In general, PCA requires to solve the following eigenvalue problem⁴:

$$\mathbf{C}\mathbf{v} = \mu\mathbf{v}. \quad (24)$$

A key observation of KPCA is that the eigenvectors lie in the span of the feature vectors, since from eq. (24) we have

$$\mathbf{v} = \frac{1}{\mu M} \sum_{i=1}^M [\phi(\mathbf{x}^i)^\top \mathbf{v}] \phi(\mathbf{x}^i) = \sum_{i=1}^M \alpha_i \phi(\mathbf{x}^i). \quad (25)$$

Here we use $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_M]^\top$ to represent the coefficients. This implies that instead of directly dealing with eq. (24) we can consider a set of n projected equations: $\phi(\mathbf{x}^i)^\top \mathbf{C}\mathbf{v} = \mu \phi(\mathbf{x}^i)^\top \mathbf{v}$, $i = 1, \dots, M$. Plugging eq. (25) here and replacing $\phi(\mathbf{x}^i)^\top \phi(\mathbf{x}^j)$ with $k(\mathbf{x}^i, \mathbf{x}^j)$ (the kernel trick), we get $\frac{1}{M} \mathbf{K}\mathbf{K}\boldsymbol{\alpha} = \mu \mathbf{K}\boldsymbol{\alpha}$, which turns out an eigenvalue problem for \mathbf{K} : $\frac{1}{M} \mathbf{K}\boldsymbol{\alpha} = \mu \boldsymbol{\alpha}$. Note that this is exactly the same eigenvalue problem solved in eq. (3). As above, we denote by $\mathbf{u}_1, \dots, \mathbf{u}_J$ the eigenvectors of \mathbf{K} that correspond to the J largest eigenvalues $\lambda_1 \geq \dots \geq \lambda_J$, and we have $\mu_j = \frac{\lambda_j}{M}$. To determine the $\boldsymbol{\alpha}$ s, we set the eigenvectors \mathbf{v} to have unit lengths: $\mathbf{v}^\top \mathbf{v} = \boldsymbol{\alpha}^\top \mathbf{K}\boldsymbol{\alpha} = \lambda \boldsymbol{\alpha}^\top \boldsymbol{\alpha} = 1$. So the

⁴We reuse some notations from the above sections (e.g., the eigenvalue μ), and as we shall see, they are closely related.

$\boldsymbol{\alpha}$ s should be normalized to have length $\frac{1}{\lambda}$: $\boldsymbol{\alpha}_j = \frac{1}{\sqrt{\lambda_j}} \mathbf{u}_j$.

For a new data point \mathbf{x} , KPCA computes the embedding $\xi(\mathbf{x})$ (the projection onto the first J eigenvectors) as $\xi(\mathbf{x}) = [\phi(\mathbf{x})^\top \mathbf{v}_1, \dots, \phi(\mathbf{x})^\top \mathbf{v}_J]^\top = [\boldsymbol{\alpha}_1^\top \mathbf{k}_\mathbf{x}, \dots, \boldsymbol{\alpha}_J^\top \mathbf{k}_\mathbf{x}]^\top$, where $\mathbf{k}_\mathbf{x} = [k(\mathbf{x}, \mathbf{x}^1), \dots, k(\mathbf{x}, \mathbf{x}^M)]^\top$. It was pointed out by Williams & Seeger (2000) to be equivalent to using the well understood Nyström approximation. We can see this by noticing that each component of $\xi(\mathbf{x})$ is identical to eq. (6) up to a scaling factor:

$$\xi_j(\mathbf{x}) = \boldsymbol{\alpha}_j^\top \mathbf{k}_\mathbf{x} = \sqrt{\frac{\lambda_j}{M}} \hat{\psi}_j(\mathbf{x}). \quad (26)$$

Looking back at eq. (16), we can see that SSGE estimates the gradients by a linear estimator with KPCA embeddings as input features. As KPCA embeddings are known to automatically adapt to the geometry of the samples, given a suitable kernel is chosen, it can reduce the curse of dimensionality when the estimator is applied to high dimensional spaces, which helps explain the effectiveness of SSGE.

Connection to Manifold-modeling Dimension Reduction Methods It has been pointed out in previous works (Williams, 2001; Ham et al., 2004; Bengio et al., 2004a;b) that many successful manifold-modeling dimension reduction methods (e.g., MDS, LLE, Laplacian eigenmaps, and Spectral clustering) can be viewed as KPCA with different ways of constructing data-dependent kernels. We believe it is a promising direction to learn a better kernel from a dataset of samples that could improve the manifold modeling behavior of KPCA embeddings, thus further improving the gradient estimator.

7. Conclusion

We propose the Spectral Stein Gradient Estimator (SSGE) for implicit distributions. Unlike previous methods, SSGE directly estimates the gradient function and thus has a principled out-of-sample extension. Future work may include learning kernels or eigenfunctions in the estimator, as indicated by the error bound as well as the connection to dimension reduction methods.

Acknowledgements

We thank anonymous reviewers for insightful feedbacks, and thank the meta-reviewer and Chang Liu for comments on improving Theorem 1. This work was supported by NSFC Projects (Nos. 61620106010, 61621136008, 61332007), Beijing NSF Project (No. L172037), Tiangong Institute for Intelligent Computing, NVIDIA NVAIL Program, Siemens and Intel.

References

- Andrieu, C., Roberts, G. O., et al. The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- Baker, C. T. *The Numerical Treatment of Integral Equations*. Clarendon Press, Oxford, 1997.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Bengio, Y., Delalleau, O., Roux, N. L., Paiement, J.-F., Vincent, P., and Ouimet, M. Learning eigenfunctions links spectral embedding and kernel PCA. *Neural computation*, 16(10):2197–2219, 2004a.
- Bengio, Y., Paiement, J.-f., Vincent, P., Delalleau, O., Roux, N. L., and Ouimet, M. Out-of-sample extensions for LLE, Isomap, MDS, eigenmaps, and spectral clustering. In *Advances in Neural Information Processing Systems*, pp. 177–184, 2004b.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks. In *International Conference on Machine Learning*, pp. 1613–1622, 2015.
- Borg, I. and Groenen, P. J. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005.
- Burges, C. J. et al. Dimension reduction: A guided tour. *Foundations and Trends® in Machine Learning*, 2(4): 275–365, 2010.
- Chwialkowski, K., Strathmann, H., and Gretton, A. A kernel test of goodness of fit. In *International Conference on Machine Learning*, pp. 2606–2615, 2016.
- Dai, B., He, N., Dai, H., and Song, L. Provable bayesian inference via particle mirror descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 985–994, 2016.
- Donahue, J., Krähenbühl, P., and Darrell, T. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., and Courville, A. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- 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 stein’s method. In *Advances in Neural Information Processing Systems*, pp. 226–234, 2015.
- Ham, J., Lee, D. D., Mika, S., and Schölkopf, B. A kernel view of the dimensionality reduction of manifolds. In *International Conference on Machine Learning*, pp. 47, 2004.
- Hernández-Lobato, J. M. and Adams, R. Probabilistic back-propagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869, 2015.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Klambauer, G., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017.
- Huszár, F. Variational inference using implicit distributions. *arXiv preprint arXiv:1702.08235*, 2017.
- Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.
- Izbicki, R., Lee, A., and Schafer, C. High-dimensional density ratio estimation with extensions to approximate likelihood computation. In *International Conference on Artificial Intelligence and Statistics*, pp. 420–429, 2014.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Li, Y. and Turner, R. E. Gradient estimators for implicit models. In *International Conference on Learning Representations*, 2018.
- Liu, Q. and Feng, Y. Two methods for wild variational inference. *arXiv preprint arXiv:1612.00081*, 2016.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances In Neural Information Processing Systems*, pp. 2370–2378, 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.
- Louizos, C. and Welling, M. Structured and efficient variational deep learning with matrix gaussian posteriors. In *International Conference on Machine Learning*, pp. 1708–1716, 2016.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *International Conference on Machine Learning*, pp. 2391–2400, 2017.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Murray, I. and Adams, R. P. Slice sampling covariance hyperparameters of latent gaussian models. In *Advances in Neural Information Processing Systems*, pp. 1732–1740, 2010.
- Neal, R. M. et al. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11), 2011.
- Nyström, E. J. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. *Acta Mathematica*, 54(1):185–204, 1930.
- Parlett, B. N. and Scott, D. S. The lanczos algorithm with selective orthogonalization. *Mathematics of computation*, 33(145):217–238, 1979.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Ranganath, R., Tran, D., Altsosaar, J., and Blei, D. Operator variational inference. In *Advances in Neural Information Processing Systems*, pp. 496–504, 2016.
- Rosasco, L., Belkin, M., and Vito, E. D. On learning with integral operators. *Journal of Machine Learning Research*, 11(Feb):905–934, 2010.
- Roweis, S. T. and Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- Schölkopf, B., Smola, A., and Müller, K.-R. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.
- Sejdinovic, D. and Gretton, A. What is an rkhs?
- Sejdinovic, D., Strathmann, H., Garcia, M. L., Andrieu, C., and Gretton, A. Kernel adaptive metropolis-hastings. In *International Conference on Machine Learning*, pp. 1665–1673, 2014.
- Shi, J., Chen, J., Zhu, J., Sun, S., Luo, Y., Gu, Y., and Zhou, Y. ZhuSuan: A library for Bayesian deep learning. *arXiv preprint arXiv:1709.05870*, 2017.
- Shi, J., Sun, S., and Zhu, J. Kernel implicit variational inference. In *International Conference on Learning Representations*, 2018.
- Sinha, K. and Belkin, M. Semi-supervised learning using sparse eigenfunction bases. In *Advances in Neural Information Processing Systems*, pp. 1687–1695, 2009.
- Sriperumbudur, B., Fukumizu, K., Gretton, A., Hyvärinen, A., and Kumar, R. Density estimation in infinite dimensional exponential families. *The Journal of Machine Learning Research*, 18(1):1830–1888, 2017.
- Stein, C. M. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, pp. 1135–1151, 1981.
- Strathmann, H., Sejdinovic, D., Livingstone, S., Szabo, Z., and Gretton, A. Gradient-free hamiltonian monte carlo with efficient kernel exponential families. In *Advances in Neural Information Processing Systems*, pp. 955–963, 2015.
- Sutherland, D., Strathmann, H., Arbel, M., and Gretton, A. Efficient and principled score estimation with nystrm kernel exponential families. In *International Conference on Artificial Intelligence and Statistics*, pp. 652–660, 2018.
- Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pp. 5529–5539, 2017.
- Weiss, Y. Segmentation using eigenvectors: a unifying view. In *The proceedings of the seventh IEEE international conference on computer vision*, volume 2, pp. 975–982. IEEE, 1999.
- Williams, C. and Seeger, M. The effect of the input density distribution on kernel-based classifiers. In *International Conference on Machine Learning*, 2000.
- Williams, C. K. On a connection between kernel PCA and metric multidimensional scaling. In *Advances in Neural Information Processing Systems*, pp. 675–681, 2001.
- Williams, C. K. and Seeger, M. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, pp. 682–688, 2001.