

---

# Learning Deep Kernels for Exponential Family Densities

---

Li K. Wenliang<sup>\*1</sup> Danica J. Sutherland<sup>\*1</sup> Heiko Strathmann<sup>1</sup> Arthur Gretton<sup>1</sup>

## Abstract

The kernel exponential family is a rich class of distributions, which can be fit efficiently and with statistical guarantees by score matching. Being required to choose *a priori* a simple kernel such as the Gaussian, however, limits its practical applicability. We provide a scheme for learning a kernel parameterized by a deep network, which can find complex location-dependent features of the local data geometry. This gives a very rich class of density models, capable of fitting complex structures on moderate-dimensional problems. Compared to deep density models fit via maximum likelihood, our approach provides a complementary set of strengths and tradeoffs: in empirical studies, deep maximum-likelihood models can yield higher likelihoods, while our approach gives better estimates of the gradient of the log density, the *score*, which describes the distribution’s shape.

## 1. Introduction

Density estimation is a foundational problem in statistics and machine learning (Devroye & Györfi, 1985; Wasserman, 2006), lying at the core of both supervised and unsupervised machine learning problems. Classical techniques such as kernel density estimation, however, struggle to exploit the structure inherent to complex problems, and thus can require unreasonably large sample sizes for adequate fits. For instance, assuming only twice-differentiable densities, the  $L_2$  risk of density estimation with  $N$  samples in  $D$  dimensions scales as  $\mathcal{O}(N^{-4/(4+D)})$  (Wasserman, 2006, Section 6.5).

One promising approach for incorporating this necessary structure is the *kernel exponential family* (Canu & Smola, 2006; Fukumizu, 2009; Sriperumbudur et al., 2017). This

---

<sup>\*</sup>Equal contribution <sup>1</sup>Gatsby Computational Neuroscience Unit, University College London, London, U.K.. Correspondence to: Li K. Wenliang <wenliang2012@gmail.com>, Danica J. Sutherland <djs@djsutherland.ml>.

model allows for any log-density which is suitably smooth under a given kernel, i.e. any function in the corresponding reproducing kernel Hilbert space. Choosing a finite-dimensional kernel recovers any classical exponential family, but when the kernel is sufficiently powerful the class becomes very rich: dense in the family of continuous probability densities on compact domains in KL, TV, Hellinger, and  $L^r$  distances (Sriperumbudur et al., 2017, Corollary 2). The normalization constant is not available in closed form, making fitting by maximum likelihood difficult, but the alternative technique of *score matching* (Hyvärinen, 2005) allows for practical usage with theoretical convergence guarantees (Sriperumbudur et al., 2017).

The choice of kernel directly corresponds to a smoothness assumption on the model, allowing one to design a kernel corresponding to prior knowledge about the target density. Yet explicitly deciding upon a kernel to incorporate that knowledge can be complicated. Indeed, previous applications of the kernel exponential family model have exclusively employed simple kernels, such as the Gaussian, with a small number of parameters (e.g. the length scale) chosen by heuristics or cross-validation (e.g. Sasaki et al., 2014; Strathmann et al., 2015). These kernels are typically *spatially invariant*, corresponding to a uniform smoothness assumption across the domain. Although such kernels are sufficient for consistency in the infinite-sample limit, the induced models can fail in practice on finite datasets, especially if the data takes differently-scaled shapes in different parts of the space. Figure 1 (left) illustrates this problem when fitting a simple mixture of Gaussians. Here there are two “correct” bandwidths, one for the broad mode and one for the narrow mode. A translation-invariant kernel must pick a single one, e.g. an average between the two, and any choice will yield a poor fit on at least part of the density.

In this work, we propose to *learn* the kernel of an exponential family directly from data. We can then achieve far more than simply tuning a length scale, instead learning location-dependent kernels that adapt to the underlying shape and smoothness of the target density. We use kernels of the form

$$k(\mathbf{x}, \mathbf{y}) = \kappa(\phi(\mathbf{x}), \phi(\mathbf{y})), \quad (1)$$

where the deep network  $\phi$  extracts features of the input and  $\kappa$  is a simple kernel (e.g. a Gaussian) on those features. These types of kernels have seen success in super-

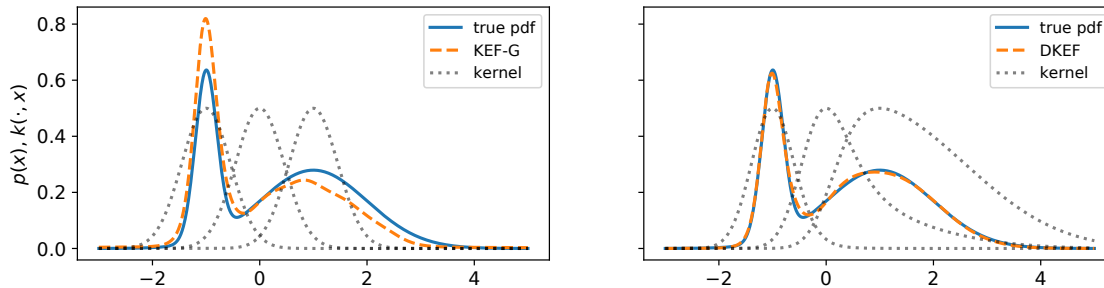


Figure 1: Fitting few samples from a Gaussian mixture, using kernel exponential families. Black dotted lines show  $k(-1, \mathbf{x})$ ,  $k(0, \mathbf{x})$ , and  $k(1, \mathbf{x})$ . (Left) Using a location-invariant Gaussian kernel, the sharper component gets too much weight. (Right) A kernel parameterized by a neural network learns length scales that adapt to the density, giving a much better fit.

vised learning (Wilson et al., 2016; Jean et al., 2018) and critic functions for implicit generative models (Li et al., 2017; Bellemare et al., 2017; Bińkowski et al., 2018; Arbel et al., 2018), among other settings. We call the resulting model a deep kernel exponential family (DKEF).

We can train both kernel parameters (including all the weights of the deep network) and, unusually, even *regularization* parameters directly on the data, in a form of meta-learning. Normally, directly optimizing regularization parameters would always yield 0, since their beneficial effect in preventing overfitting is by definition not seen on the training set. Here, though, we can exploit the closed-form fit of the kernel exponential family to optimize a “held-out” score (Section 3). Figure 1 (right) demonstrates the success of this model on the same mixture of Gaussians; here the learned, location-dependent kernel gives a much better fit.

We compare the results of our new model to recent general-purpose deep density estimators, primarily *autoregressive models* (Uribe et al., 2013; Germain et al., 2015; van den Oord et al., 2016) and *normalizing flows* (Jimenez Rezende & Mohamed, 2015; Dinh et al., 2017; Papamakarios et al., 2017). These models learn deep networks with structures designed to compute normalized densities, and are fit via maximum likelihood. We explore the strengths and limitations of both deep likelihood models and deep kernel exponential families on a variety of datasets, including artificial data designed to illustrate scenarios where certain surprising problems arise, as well as benchmark datasets used previously in the literature. The models fit by maximum likelihood typically give somewhat higher likelihoods, whereas the deep kernel exponential family generally better fits the *shape* of the distribution.

## 2. Background

**Score matching** Suppose we observe  $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$ , a set of independent samples  $\mathbf{x}_n \in \mathbb{R}^D$  from an unknown density  $p_0(\mathbf{x})$ . We posit a class of possible models  $\{p_\theta\}$ , pa-

rameterized by  $\theta$ ; our goal is to use the data  $\mathcal{D}$  to select some  $\hat{\theta}$  such that  $p_{\hat{\theta}} \approx p_0$ . The standard approach for selecting  $\theta$  is maximum likelihood:  $\hat{\theta} = \arg \max_{\theta} \prod_{n=1}^N p_\theta(\mathbf{x}_n)$ .

Many interesting model classes, however, are defined as  $p_\theta(\mathbf{x}) = \tilde{p}_\theta(\mathbf{x})/Z_\theta$ , where the normalization constant  $Z_\theta = \int_{\mathbf{x}} \tilde{p}_\theta(\mathbf{x}) d\mathbf{x}$  cannot be easily computed. In this setting, an optimization algorithm to estimate  $\theta$  by maximum likelihood requires estimating (the derivative of)  $Z_\theta$  for each candidate  $\theta$  considered during optimization. Moreover, the maximum likelihood solution may not even be well-defined when  $\theta$  is infinite-dimensional (Barron & Sheu, 1991; Fukumizu, 2009). The intractability of maximum likelihood led Hyvärinen (2005) to propose an alternative objective, called *score matching*. Rather than maximizing the likelihood, one minimizes the Fisher divergence  $J(p_\theta \| p_0)$ :

$$\frac{1}{2} \int p_0(\mathbf{x}) \|\nabla_{\mathbf{x}} \log p_\theta(\mathbf{x}) - \nabla_{\mathbf{x}} \log p_0(\mathbf{x})\|_2^2 d\mathbf{x}. \quad (2)$$

Under mild regularity conditions, this is equal to

$$\int_{\mathbf{x}} p_0(\mathbf{x}) \sum_{d=1}^D \left[ \partial_d^2 \log p_\theta(\mathbf{x}) + \frac{1}{2} (\partial_d \log p_\theta(\mathbf{x}))^2 \right] d\mathbf{x}, \quad (3)$$

up to an additive constant depending only on  $p_0$ , which can be ignored during training. Here  $\partial_d^n f(\mathbf{x})$  denotes  $\frac{\partial^n}{(\partial y_d)^n} f(\mathbf{y})|_{\mathbf{y}=\mathbf{x}}$ . We can estimate (3) with  $\hat{J}(p_\theta, \mathcal{D})$ :

$$\frac{1}{N} \sum_{n=1}^N \sum_{d=1}^D \left[ \partial_d^2 \log p_\theta(\mathbf{x}_n) + \frac{1}{2} (\partial_d \log p_\theta(\mathbf{x}_n))^2 \right]. \quad (4)$$

Notably, (4) does not depend on  $Z_\theta$ , and so we can minimize it to find an unnormalized model  $\tilde{p}_{\hat{\theta}}$  for  $p_0$ . Score matching is consistent in the well-specified setting (Hyvärinen, 2005, Theorem 2), and was related to maximum likelihood by Lyu (2009), who argues it finds a fit robust to infinitesimal noise.

Unnormalized models  $\tilde{p}$  are sufficient for many tasks (LeCun et al., 2006), including finding modes, approximating Hamiltonian Monte Carlo on targets without gradients

(Strathmann et al., 2015), and learning discriminative features (Janzamin et al., 2014). If we require a normalized model, however, we can estimate the normalizing constant once, after estimating  $\theta$ ; this will be far more computationally efficient than estimating it at each step of an iterative maximum likelihood optimization algorithm.

**Kernel exponential families** The kernel exponential family (Canu & Smola, 2006; Sriperumbudur et al., 2017) is the class of all densities satisfying a smoothness constraint:  $\log \tilde{p}(\mathbf{x}) = f(\mathbf{x}) + \log q_0(\mathbf{x})$ , where  $q_0$  is some fixed function and  $f$  is any function in the reproducing kernel Hilbert space  $\mathcal{H}$  with kernel  $k$ . This class is an exponential family with natural parameter  $f$  and sufficient statistic  $k(\mathbf{x}, \cdot)$ , due to the reproducing property  $f(\mathbf{x}) = \langle f, k(\mathbf{x}, \cdot) \rangle_{\mathcal{H}}$ :

$$\tilde{p}_f(\mathbf{x}) = \exp(f(\mathbf{x})) q_0(\mathbf{x}) = \exp(\langle f, k(\mathbf{x}, \cdot) \rangle_{\mathcal{H}}) q_0(\mathbf{x}).$$

Using a simple finite-dimensional  $\mathcal{H}$ , we can recover any standard exponential family, e.g. normal, gamma, or Poisson; if  $\mathcal{H}$  is sufficiently rich, the family can approximate any continuous distribution with tails like  $q_0$  arbitrarily well (Sriperumbudur et al., 2017, Example 1 and Corollary 2).

These models do not in general have a closed-form normalizer. For some  $f$  and  $q_0$ ,  $\tilde{p}_f$  may not even be normalizable, but if  $q_0$  is e.g. a Gaussian density, typical choices of  $\phi$  and  $\kappa$  in (1) guarantee a normalizer exists (Appendix A).

Sriperumbudur et al. (2017) proved good statistical properties for choosing  $f \in \mathcal{H}$  by minimizing a regularized form of (4),  $\hat{f} = \arg \min_{f \in \mathcal{H}} \hat{J}(\tilde{p}_f, \mathcal{D}) + \lambda \|f\|_{\mathcal{H}}^2$ , but their algorithm has an impractical computational cost of  $\mathcal{O}(N^3 D^3)$ . This can be alleviated with the Nyström-type “lite” approximation (Strathmann et al., 2015; Sutherland et al., 2018): select  $M$  inducing points  $\mathbf{z}_m \in \mathbb{R}^D$ , and select  $f \in \mathcal{H}$  as

$$f_{\alpha, \mathbf{z}}^k(\mathbf{x}) = \sum_{m=1}^M \alpha_m k(\mathbf{x}, \mathbf{z}_m), \quad \tilde{p}_{\alpha, \mathbf{z}}^k = \tilde{p}_{f_{\alpha, \mathbf{z}}^k}. \quad (5)$$

As the span of  $\{k(\mathbf{z}, \cdot)\}_{\mathbf{z} \in \mathbb{R}^D}$  is dense in  $\mathcal{H}$ , this is a natural approximation, similar to classical RBF networks (Broomhead & Lowe, 1988). The “lite” model often yields excellent empirical results at much lower computational cost than the full estimator. We can regularize (4) in several ways and still find a closed-form solution for  $\alpha$ . In this work, our loss  $\hat{J}(f_{\alpha, \mathbf{z}}^k, \lambda, \mathcal{D})$  will be

$$\hat{J}(\tilde{p}_{\alpha, \mathbf{z}}^k, \mathcal{D}) + \frac{\lambda_{\alpha}}{2} \|\alpha\|^2 + \frac{\lambda_C}{2N} \sum_{n=1}^N \sum_{d=1}^D [\partial_d^2 \log \tilde{p}_{\alpha, \mathbf{z}}^k(\mathbf{x}_n)]^2.$$

Sutherland et al. (2018) used a small  $\lambda_{\alpha}$  for numerical stability but primarily regularized with  $\lambda_H \|f_{\alpha, \mathbf{z}}^k\|_{\mathcal{H}}^2$ . As we change  $k$ , however,  $\|f\|_{\mathcal{H}}$  changes meaning, and we found empirically that this regularizer tends to harm the fit. The

$\lambda_C$  term was recommended by Kingma & LeCun (2010), encouraging the learned log-density to be smooth without much extra computation; it provides some empirical benefit in our context. Given  $k$ ,  $\mathbf{z}$ , and  $\lambda$ , Proposition 3 (Appendix B) shows we can find the optimal  $\alpha$  by solving an  $M \times M$  linear system in  $\mathcal{O}(M^2 ND + M^3)$  time: the  $\alpha$  which minimizes  $\hat{J}(f_{\alpha, \mathbf{z}}^k, \lambda, \mathcal{D})$  is

$$\alpha(\lambda, k, \mathbf{z}, \mathcal{D}) = -(\mathbf{G} + \lambda_{\alpha} \mathbf{I} + \lambda_C \mathbf{U})^{-1} \mathbf{b} \quad (6)$$

$$G_{m, m'} = \frac{1}{N} \sum_{n=1}^N \sum_{d=1}^D \partial_d k(\mathbf{x}_n, \mathbf{z}_m) \partial_d k(\mathbf{x}_n, \mathbf{z}_{m'})$$

$$U_{m, m'} = \frac{1}{N} \sum_{n=1}^N \sum_{d=1}^D \partial_d^2 k(\mathbf{x}_n, \mathbf{z}_m) \partial_d^2 k(\mathbf{x}_n, \mathbf{z}_{m'})$$

$$b_m = \frac{1}{N} \sum_{n=1}^N \sum_{d=1}^D \partial_d^2 k(\mathbf{x}_n, \mathbf{z}_m) + \partial_d \log q_0(\mathbf{x}_n) \partial_d k(\mathbf{x}_n, \mathbf{z}_m) + \lambda_C \partial_d^2 \log q_0(\mathbf{x}_n) \partial_d^2 k(\mathbf{x}_n, \mathbf{z}_m).$$

### 3. Fitting Deep Kernels

All previous applications of score matching in the kernel exponential family of which we are aware (e.g. Sriperumbudur et al., 2017; Strathmann et al., 2015; Sun et al., 2015; Sutherland et al., 2018) have used kernels of the form  $k(\mathbf{x}, \mathbf{y}) = \exp(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|^2) + r(\mathbf{x}^T \mathbf{y} + c)^2$ , with kernel parameters and regularization weights either fixed a priori or selected via cross-validation. This simple form allows the various kernel derivatives required in (6) to be easily computed by hand, and the small number of parameters makes grid search adequate for model selection. But, as discussed in Section 1, these simple kernels are insufficient for complex datasets. Thus we wish to use a richer class of kernels  $\{k_{\mathbf{w}}\}$ , with a large number of parameters  $\mathbf{w}$  – in particular, kernels defined by a neural network. This prohibits model selection via simple grid search.

One could attempt to directly minimize  $\hat{J}(f_{\alpha, \mathbf{z}}^{k_{\mathbf{w}}}, \lambda, \mathcal{D})$  jointly in the kernel parameters  $\mathbf{w}$ , the model parameters  $\alpha$ , and perhaps the inducing points  $\mathbf{z}$ . Consider, however, the case where we simply use a Gaussian kernel and  $\{\mathbf{z}_m\} = \mathcal{D}$ . Then we can achieve arbitrarily good values of (3) by taking  $\sigma \rightarrow 0$ , drastically overfitting to the training set  $\mathcal{D}$ .

We can avoid this problem – and additionally find the best values for the regularization weights  $\lambda$  – with a form of meta-learning. We find choices for the kernel and regularization which will give us a good value of  $\hat{J}$  on a “validation set”  $\mathcal{D}_v$  when fit to a fresh “training set”  $\mathcal{D}_t$ . Specifically, we take stochastic gradient steps following  $\nabla_{\lambda, \mathbf{w}, \mathbf{z}} \hat{J}(\tilde{p}_{\alpha(\lambda, k_{\mathbf{w}}, \mathbf{z}, \mathcal{D}_t), \mathbf{z}}^{k_{\mathbf{w}}}, \mathcal{D}_v)$ . We can easily do this because we have a differentiable closed-form expression (6) for the fit to  $\mathcal{D}_t$ , rather than having to e.g. back-propagate through an unrolled iterative optimization procedure. As we

---

**Algorithm 1:** Full training procedure
 

---

**input:** Dataset  $\mathcal{D}$ ; initial inducing points  $\mathbf{z}$ , kernel parameters  $\mathbf{w}$ , regularization  $\boldsymbol{\lambda} = (\lambda_\alpha, \lambda_C)$

Split  $\mathcal{D}$  into  $\mathcal{D}_1$  and  $\mathcal{D}_2$ ;

Optimize  $\mathbf{w}$ ,  $\boldsymbol{\lambda}$ ,  $\mathbf{z}$ , and maybe  $q_0$  params:

**while**  $\hat{J}(\tilde{p}_{\alpha}^{k_w}(\boldsymbol{\lambda}, k_w, \mathbf{z}, \mathcal{D}_1), \mathcal{D}_2)$  still improving **do**

    Sample disjoint data subsets  $\mathcal{D}_t, \mathcal{D}_v \subset \mathcal{D}_1$ ;

$f(\cdot) = \sum_{m=1}^M \alpha_m(\boldsymbol{\lambda}, k_w, \mathbf{z}, \mathcal{D}_t) k_w(\mathbf{z}_m, \cdot)$ ;

$\hat{J} = \frac{1}{|\mathcal{D}_v|} \sum_{n=1}^{|\mathcal{D}_v|} \sum_{d=1}^D [\partial_d^2 f(\mathbf{x}_n) + \frac{1}{2}(\partial_d f(\mathbf{x}_n))^2]$ ;

    Take SGD step in  $\hat{J}$  for  $\mathbf{w}$ ,  $\boldsymbol{\lambda}$ ,  $\mathbf{z}$ , maybe  $q_0$  params;

**end**

Optimize  $\boldsymbol{\lambda}$  for fitting on larger batches:

**while**  $\hat{J}(\tilde{p}_{\alpha}^{k_w}(\boldsymbol{\lambda}, k_w, \mathbf{z}, \mathcal{D}_1), \mathcal{D}_2)$  still improving **do**

$f(\cdot) = \sum_{m=1}^M \alpha_m(\boldsymbol{\lambda}, k_w, \mathbf{z}, \mathcal{D}_1) k_w(\cdot, \mathbf{z}_m)$ ;

    Sample subset  $\mathcal{D}_v \subset \mathcal{D}_2$ ;

$\hat{J} = \frac{1}{|\mathcal{D}_v|} \sum_{n=1}^{|\mathcal{D}_v|} \sum_{d=1}^D [\partial_d^2 f(\mathbf{x}_n) + \frac{1}{2}(\partial_d f(\mathbf{x}_n))^2]$ ;

    Take SGD steps in  $\hat{J}$  for  $\boldsymbol{\lambda}$  only;

**end**

Finalize  $\alpha$  on  $\mathcal{D}_1$ :

Find  $\alpha = \alpha(\boldsymbol{\lambda}, k_w, \mathbf{z}, \mathcal{D}_1)$ ;

**return:**  $\log \tilde{p}(\cdot) = \sum_{m=1}^M \alpha_m k_w(\cdot, \mathbf{z}_m) + \log q_0(\cdot)$ ;

---

used small minibatches in this procedure, for the final fit we use the whole dataset: we first freeze  $\mathbf{w}$  and  $\mathbf{z}$  and find the optimal  $\boldsymbol{\lambda}$  for the whole training data, then finally fit  $\alpha$  with the new  $\boldsymbol{\lambda}$ . This process is summarized in Algorithm 1.

**Computing kernel derivatives** Solving for  $\alpha$  and computing the loss (4) require matrices of kernel second derivatives, but current deep learning-oriented automatic differentiation systems are not optimized for evaluating tensor-valued higher-order derivatives at once. We therefore implement backpropagation to compute  $\mathbf{G}$ ,  $\mathbf{U}$ , and  $\mathbf{b}$  of (6) as TensorFlow operations (Abadi et al., 2015) to obtain the scalar loss  $\hat{J}$ , and used TensorFlow’s automatic differentiation only to optimize  $\mathbf{w}$ ,  $\mathbf{z}$ ,  $\boldsymbol{\lambda}$ , and  $q_0$  parameters.

Backpropagation to find these second derivatives requires explicitly computing the Hessians of intermediate layers of the network, which becomes quite expensive as the model grows; this limits the size of kernels that our model can use. A more efficient implementation based on Hessian-vector products might be possible in an automatic differentiation system with better support for matrix derivatives.

**Kernel architecture** We will choose our kernel  $k_w(\mathbf{x}, \mathbf{y})$  as a mixture of  $R$  Gaussian kernels with length scales  $\sigma_r$ , taking in features of the data extracted by a network  $\phi_{w_r}(\cdot)$ :

$$\sum_{r=1}^R \rho_r \exp\left(-\frac{1}{2\sigma_r^2} \|\phi_{w_r}(\mathbf{x}) - \phi_{w_r}(\mathbf{y})\|^2\right). \quad (7)$$

Combining  $R$  components makes it easier to account for both short-range and long-range dependencies. We constrain  $\rho_r \geq 0$  to ensure a valid kernel, and  $\sum_{r=1}^R \rho_r = 1$  for simplicity. The networks  $\phi_w$  are made of  $L$  fully connected layers of width  $W$ . For  $L > 1$ , we found that adding a skip connection from data directly to the top layer speeds up learning. A softplus nonlinearity,  $\log(1 + \exp(x))$ , ensures that the model is twice-differentiable so (3) is well-defined.

### 3.1. Behavior on Mixtures

One interesting limitation of score matching is the following: suppose that  $p_0$  is composed of two disconnected components,  $p_0(\mathbf{x}) = \pi p_1(\mathbf{x}) + (1 - \pi)p_2(\mathbf{x})$  for  $\pi \in (0, 1)$  and  $p_1, p_2$  having disjoint, separated support. Then  $\nabla \log p_0(\mathbf{x})$  will be  $\nabla \log p_1(\mathbf{x})$  in the support of  $p_1$ , and  $\nabla \log p_2(\mathbf{x})$  in the support of  $p_2$ . Score matching compares  $\nabla \log \tilde{p}_\theta$  to  $\nabla \log p_1$  and  $\nabla \log p_2$ , but is completely blind to  $\tilde{p}_\theta$ ’s relative mass between the two components; it is equally happy with *any* reweighting of the components.

If all modes are connected by regions of positive density, then the log density gradient in between components will determine their relative weight, and indeed score matching is then consistent. But when  $p_0$  is *nearly* zero between two dense components, so that there are no or few samples in between, score matching will generally have insufficient evidence to weight nearly-separate components.

Proposition 4 (Appendix C) studies the the kernel exponential family in this case. For two components that are completely separated according to  $k$ , (6) fits each as it would if given only that component, except that the effective  $\lambda_\alpha$  is scaled: smaller components are regularized more.

Appendix C.1 studies a simplified case where the kernel length scale  $\sigma$  is far wider than the component; then the density ratio between components, which should be  $\frac{\pi}{1-\pi}$ , is approximately  $\exp\left(\frac{D}{2\sigma^2\lambda_\alpha} \left(\pi - \frac{1}{2}\right)\right)$ . Depending on the value of  $\frac{D}{2\sigma^2\lambda_\alpha}$ , this ratio will often either be quite extreme, or nearly 1. It is unclear, however, how well this result generalizes to other settings.

A heuristic workaround when disjoint components are suspected is as follows: run a clustering algorithm to identify disjoint components, separately fit a model to each cluster, then weight each model according to its sample count. When the components are well-separated, this clustering is straightforward, but it may be difficult in high-dimensional cases when samples are sparse but not fully separated.

### 3.2. Model Evaluation

In addition to qualitatively evaluating fits, we will evaluate our models with three quantitative criteria. The first is the finite-set Stein discrepancy (FSSD; Jitkrittum et al.,

2017), a measure of model fit which does not depend on the normalizer  $Z_\theta$ . It examines the fit of the model at  $J$  test locations  $\mathbf{V} = \{\mathbf{v}_b\}_{b=1}^B$  using an kernel  $l(\cdot, \cdot)$ , as  $\frac{1}{DB} \sum_{b=1}^B \|\mathbb{E}_{\mathbf{x} \sim p_0} [l(\mathbf{x}, \mathbf{v}_b) \nabla_{\mathbf{x}} \log p(\mathbf{x}) + \nabla_{\mathbf{x}} l(\mathbf{x}, \mathbf{v}_b)]\|^2$ . With randomly selected  $\mathbf{V}$  and some mild assumptions, it is zero if and only if  $p = p_0$ . We use a Gaussian kernel with bandwidth equal to the median distance between test points,<sup>1</sup> and choose  $\mathbf{V}$  by adding small Gaussian noise to data points. Jitkrittum et al. (2018) construct a hypothesis test to test which of  $p$  and  $p'$  is closer to  $p_0$  in the FSSD. We will report a score – the  $p$ -value of this test – which is near 0 when model  $p$  is better, near 1 when model  $p'$  is better, and around  $\frac{1}{2}$  when the two models are equivalent. We emphasize that we are using this as a model comparison score on an interpretable scale, but *not* following a hypothesis testing framework. Another similar performance measure is the kernel Stein discrepancy (KSD) (Chwialkowski et al., 2016), where the model’s goodness-of-fit is evaluated at all test data points rather than at random test locations. We omit the results as they are essentially identical to that of the FSSD, even across a wide range of kernel bandwidths.

As all our models are twice-differentiable, we also compare the score-matching loss (4) on held-out test data. A lower score-matching loss implies a smaller Fisher divergence between model and data distributions.

Finally, we compare test log-likelihoods, using importance sampling estimates of the normalizing constant  $Z_\theta$ :

$$\hat{Z}_\theta = \frac{1}{U} \sum_{u=1}^U r_u \quad \text{where } r_u := \frac{\tilde{p}_\theta(\mathbf{y}_u)}{q_0(\mathbf{y}_u)}, \quad \mathbf{y}_u \sim q_0,$$

so  $\mathbb{E} \hat{Z}_\theta = \int \frac{\tilde{p}_\theta(\mathbf{y}_u)}{q_0(\mathbf{y}_u)} q_0(\mathbf{y}_u) = Z_\theta$ . Our log-likelihood estimate is  $\log \hat{p}_\theta(\mathbf{x}) = \log \tilde{p}_\theta(\mathbf{x}) - \log \hat{Z}_\theta$ . This estimator is consistent, but Jensen’s inequality tells us that  $\mathbb{E} \log \hat{p}_\theta(\mathbf{x}) > \log p_\theta(\mathbf{x})$ , so our evaluation will be over-optimistic. Worse, the variance of  $\log \hat{Z}_\theta$  can be misleadingly small when the bias is still quite large; we observed this in our experiments. We can, though, bound the bias:

**Proposition 1.** *Suppose that  $a, s \in \mathbb{R}$  are such that  $\Pr(r_u \geq a) = 1$  and  $\Pr(r_u \leq s) \leq \rho < \frac{1}{2}$ . Define  $t := (s + a)/2$ ,  $\psi(q, Z_\theta) := \log \frac{Z}{q} + \frac{q}{Z} - 1$ , and let  $P := \max(\psi(a, Z_\theta), \psi(t, Z_\theta))$ . Then*

$$\log Z_\theta - \mathbb{E} \log \hat{Z}_\theta \leq \frac{\psi(t, Z_\theta)}{(Z_\theta - t)^2} \frac{\text{Var}[r_u]}{U} + P(4\rho(1 - \rho))^{\frac{U}{2}}.$$

(Proof in Appendix D.) We can find  $a, s$  because we propose from  $q_0$ , and thus we can effectively estimate the bound

<sup>1</sup>This reasonable choice avoids tuning any parameters. We do not optimize the kernel or the test locations to avoid a situation in which model  $p$  is better than  $p'$  in some respects but  $p'$  better than  $p$  in others; instead, we use a simple default mode of comparison.

(Appendix D.1). This estimate of the upper bound is itself biased upwards (Proposition 6), so it is likely, though not guaranteed, that the estimate overstates the amount of bias.

### 3.3. Previous Attempts at Deep Score Matching

Kingma & LeCun (2010) used score matching to train a (one-layer) network to output an unnormalized log-density. This approach is essentially a special case of ours: use the kernel  $k_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \phi_{\mathbf{w}}(\mathbf{x})\phi_{\mathbf{w}}(\mathbf{y})$ , where  $\phi_{\mathbf{w}} : \mathbb{R}^D \rightarrow \mathbb{R}$ . Then the function  $f_{\alpha, z}^{k_{\mathbf{w}}}(\mathbf{x})$  from (5) is

$$\sum_{m=1}^M \alpha_m \phi_{\mathbf{w}}(\mathbf{z}_m) \phi_{\mathbf{w}}(\mathbf{x}) = \left[ \sum_{m=1}^M \alpha_m \phi_{\mathbf{w}}(\mathbf{z}_m) \right] \phi_{\mathbf{w}}(\mathbf{x}).$$

The scalar in brackets is fit analytically, so  $\log p$  is determined almost entirely by the network  $\phi_{\mathbf{w}}$  plus  $\log q_0(\mathbf{x})$ .

Saremi et al. (2018) recently also attempted parameterizing the unnormalizing log-density as a deep network, using an approximation called Parzen score matching. This approximation requires a global constant bandwidth to define the Parzen window size for the loss function, fit to the dataset before learning the model. This is likely only sensible on datasets for which simple fixed-bandwidth kernel density estimation is appropriate; on more complex datasets, the loss may be poorly motivated. It also leads to substantial over-smoothing visible in their results. As they did not provide code for their method, we do not compare to it empirically.

## 4. Experimental Results

In our experiments, we compare to several alternative methods. The first group are all fit by maximum likelihood, and broadly fall into (at least) one of two categories: *autoregressive models* decompose  $p(x_1, \dots, x_D) = \prod_{d=1}^D p(x_d | \mathbf{x}_{\leq d})$  and learn a parametric density model for each of these conditionals. *Normalizing flows* instead apply a series of invertible transformations to some simple initial density, say standard normal, and then compute the density of the overall model via the Jacobian of the transformation. We use implementations<sup>2</sup> of the following several models in these categories by Papamakarios et al. (2017):

**MADE** (Germain et al., 2015) masks the connections of an autoencoder so it is autoregressive. We use two hidden layers, and each conditional a Gaussian. **MADE-MOG** is the same but with each conditional a mixture of 10 Gaussians.

**Real NVP** (Dinh et al., 2017) is a normalizing flow; we use a general-purpose form for non-image datasets.

**MAF** (Papamakarios et al., 2017). A combination of a normalizing flow and MADE, where the base density is modeled by MADE, with 5 autoregressive layers. **MAF-**

<sup>2</sup>[github.com/gpapamak/maf](https://github.com/gpapamak/maf)

**MOG** instead models the base density by MADE-MOG.

For the models above, we use layers of width 30 for experiments on synthetic data, and 100 for benchmark datasets. Larger values did not improve performance.

**KCEF** (Arbel & Gretton, 2018). Inspired by autoregressive models, the density is modeled by a cascade of kernel conditional exponential family distributions, fit by score matching with Gaussian kernels.<sup>3</sup>

**DKEF**. On synthetic datasets, we consider four variants of our model with one kernel component,  $R = 1$ . KEF-G refers to the model using a Gaussian kernel with a learned bandwidth. DKEF-G-15 has the kernel (7), with  $L = 3$  layers of width  $W = 15$ . DKEF-G-50 is the same with  $W = 50$ . To investigate whether the top Gaussian kernel helps performance, we also train DKEF-L-50, whose kernel is  $k_\theta(\mathbf{x}, \mathbf{y}) = \phi_w(\mathbf{x}) \cdot \phi_w(\mathbf{y})$ , where  $\phi_w$  has  $W = 50$ . To compare with the architecture of Kingma & LeCun (2010), DKEF-L-50-1 has the same architecture as DKEF-L-50 except that we add an extra layer with a single neuron, and use  $M = 1$ . In all experiments,  $q_0(\mathbf{x}) = \prod_{d=1}^D \exp(-|x_d - \mu_d|^{\beta_d} / (2\sigma_d^2))$ , with  $\beta_d > 1$ . On benchmark datasets, we use DKEF-G-50 and KEF-G with three kernel components,  $R = 3$ . Code for DKEF is at [github.com/kevin-w-li/deep-kexpfam](https://github.com/kevin-w-li/deep-kexpfam).

#### 4.1. Behavior on Synthetic Datasets

We first demonstrate the behavior of the models on several two-dimensional synthetic datasets Funnel, Banana, Ring, Square, Cosine, Mixture of Gaussians (MoG) and Mixture of Rings (MoR). Together, they cover a range of geometric complexities and multimodality.

We visualize the fit of various methods by showing the log density function in Figure 2. For each model fit on each distribution, we report the normalized log likelihood (left) and Fisher divergence (right). In general, the kernel score matching methods find cleaner boundaries of the distributions, and our main model DKEF-G-15 produces the lowest Fisher divergence on many of the synthetic datasets while maintaining high likelihoods.

Among the kernel exponential families, DKEF-G outperformed previous versions where ordinary Gaussian kernels are used for either joint (KEF-G) or autoregressive (KCEF) modeling. DKEF-G-50 does not substantially improve over DKEF-G-15; we omit it for space. We can gain additional insights into the model by looking at the shape of the learned kernel, shown by the colored lines; the kernels do indeed adapt to the local geometry.

DKEF-L-50 and DKEF-L-50-1 show good performance when the target density has simple geometries, but had trou-

ble in more complex cases, even with much larger networks than used by DKEF-G-15. It seems that a Gaussian kernel with inducing points provides much stronger representational features than a using linear kernel and/or a single inducing point. A large enough network  $\phi_w$  would likely be able to perform the task well, but, using currently available software, the second derivatives in the score matching loss limit our ability to use very large networks. A similar phenomenon was observed by Bińkowski et al. (2018) in the context of GAN critics, where combining some analytical RKHS optimization with deep networks allowed much smaller networks to work well.

As expected, models fit by DKEFs generally have smaller Fisher divergences than likelihood-based methods. For Funnel and Banana, the true densities are simple transformations of Gaussians, and the normalizing flow models perform relatively well. But on Ring, Square, and Cosine, the shape of the learned distribution by likelihood-based methods exhibits noticeable artifacts, especially at the boundaries. These artifacts, particularly the “breaks” in Ring, may be caused by a normalizing flow’s need to be invertible and smooth. The shape learned by DKEF-G-15 is much cleaner.

On multimodal distributions with disjoint components, likelihood-based and score matching-based methods show interesting failure modes. The likelihood-based models often connect separated modes with a “bridge”, even for MADE-MOG and MAF-MOG which use explicit mixtures. On the other hand, DKEF is able to find the shapes of the components, but the weighting between modes is unstable. As suggested in Section 3.1, we also fit mixtures of all models (except KCEF) on a partition of MoR found by spectral clustering (Shi & Malik, 2000); DKEF-G-15 produced an excellent fit.

Another challenge we observed in our experiments is that the estimator of the objective function,  $\hat{J}$ , tends to be more noisy as the model fit improves. This happens particularly on datasets where there are “sharp” features, such as Square (see Figure 4 in Appendix E.1), where the model’s curvature becomes extreme at some points. This can cause higher variance in the gradients of the parameters, and more difficulty in optimization.

#### 4.2. Results on Benchmark Datasets

Following recent work on density estimation (Uria et al., 2013; Arbel & Gretton, 2018; Papamakarios et al., 2017; Germain et al., 2015), we trained DKEF and the likelihood-based models on five UCI datasets (Dheeru & Karra Taniskidou, 2017); in particular, we used RedWine, WhiteWine, Parkinson, HepMass, and MiniBoone. All performances were measured on held-out test sets. We did not run KCEF due to its computational expense. Appendix E.2 gives further details.

<sup>3</sup>[github.com/MichaelArbel/KCEF](https://github.com/MichaelArbel/KCEF)




Figure 2: Log densities learned by different models. Our model is DKEF-G-15 at the bottom row. Columns are different synthetic datasets. The rightmost columns shows a mixture of each model (except KCEF) on the same clustering of MoR. We subtracted the maximum from each log density, and clipped the minimum value at  $-9$ . Above each panel are shown the average log-likelihoods (left) and Fisher divergence (right) on held-out data points. Bold indicates the best fit. For DKEF-G models, faint colored lines correspond to contours at 0.9 of the kernel evaluated at different locations.

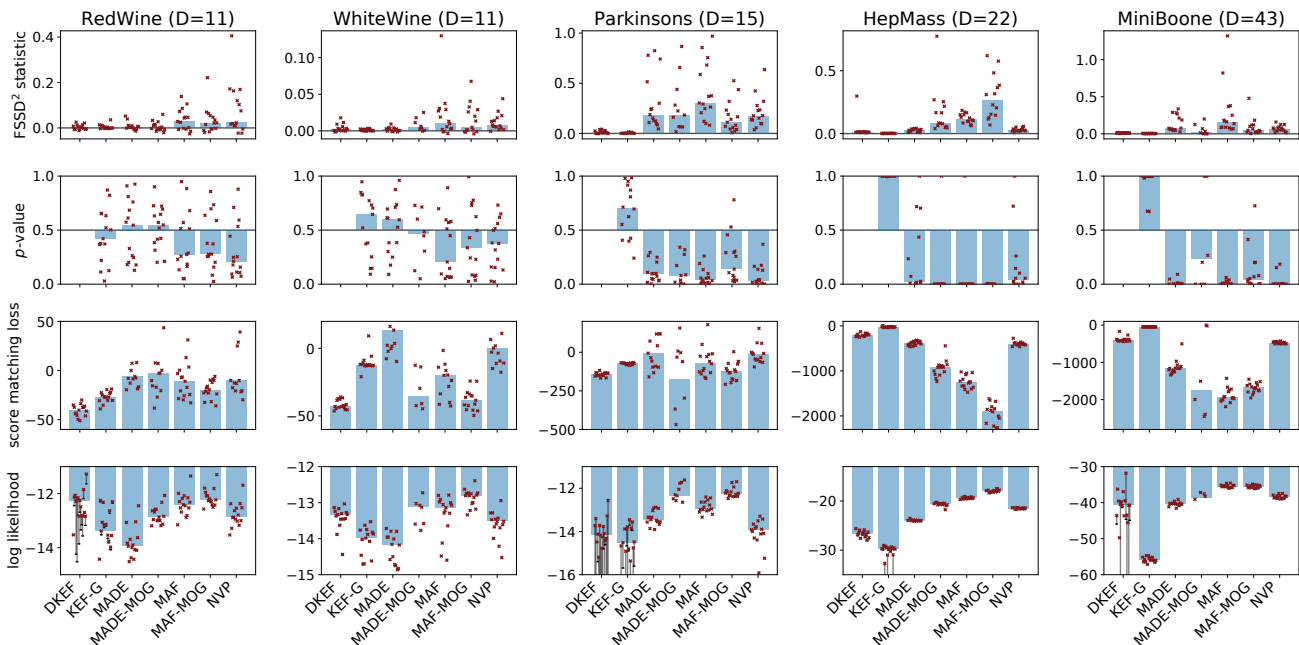


Figure 3: Results on the real datasets; bars show medians, points show each of 15 individual runs, excluding invalid values. (1st row) The estimate of the squared FSSD, a measure of model goodness of fit based on derivatives of the log density; lower is better. (2nd row) The  $p$ -value of a test that each model is no better than DKEF in terms of the FSSD; values near 0 indicate that DKEF fits the data significantly better than the other model. (3rd row) Value of the loss (4); lower is better. (4th row) Log-likelihoods; higher is better. DKEF estimates are based on  $10^{10}$  samples for  $\hat{\theta}$ , with vertical lines showing the upper bound on the bias from Proposition 1 (which is often too small to be easily visible).

Figure 3 shows results. In gradient matching as measured by the FSSD, DKEF tends to have the best values. Test set sizes are too small to yield a confident  $p$ -value on the Wine datasets, but the model comparison test confidently favors DKEF on datasets with large-enough test sets. The FSSD results agree with KSD, which is omitted. In the score matching loss<sup>4</sup> (3), DKEF is the best on Wine datasets and most runs on Parkinson, but worse on Hepmass and MiniBoone. FSSD is a somewhat more “global” measure of shape, and is perhaps more weighted towards the bulk of the distribution rather than the tails.<sup>5</sup> In likelihoods, DKEF is comparable to other methods except MADE on Wines but worse on the other, larger, datasets. Note that we trained DKEF while adding Gaussian noise with standard deviation 0.05 to the (whitened) dataset; training without noise improves the score matching loss but harms likelihood, while producing similar results for FSSD.

Results for models with a fixed Gaussian  $q_0$  were similar (Figure 6, in Appendix E.2).

<sup>4</sup>The implementation of Papamakarios et al. (2017) sometimes produced NaN values for the required second derivatives, especially for MADE-MOG. We discarded those runs for these plots.

<sup>5</sup>With a kernel approaching a Dirac delta, the FSSD<sup>2</sup> is similar to  $\text{KSD}^2 \approx \int p_0(\mathbf{x})^2 \|\nabla \log p(\mathbf{x}) - \nabla \log \tilde{p}_\theta(\mathbf{x})\|^2 d\mathbf{x}$ ; compare to  $J = \frac{1}{2} \int p_0(\mathbf{x}) \|\nabla \log p(\mathbf{x}) - \nabla \log \tilde{p}_\theta(\mathbf{x})\|^2 d\mathbf{x}$ .

## 5. Discussion

Learning deep kernels helps make the kernel exponential family practical for large, complex datasets of moderate dimension. We can exploit the closed-form fit of the  $\alpha$  vector to optimize kernel and even regularization parameters using a “held-out” loss, in a particularly convenient instance of meta-learning. We are thus able to find smoothness assumptions that fit our particular data, rather than arbitrarily choosing them a priori.

Computational expense makes score matching with deep kernels difficult to scale to models with large kernel networks, limiting the dimensionality of possible targets. Combining with the kernel conditional exponential family might help alleviate that problem by splitting the model up into several separate but marginally complex components. The kernel exponential family, and score matching in general, also struggles to correctly balance datasets with nearly-disjoint components, but it seems to generally learn density *shapes* better than maximum likelihood-based deep approaches.

## ACKNOWLEDGMENTS

This work was supported by the Gatsby Charitable Foundation. We thank Heishiro Kanagawa for helpful discussions.



## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- Arbel, M. and Gretton, A. Kernel conditional exponential family. In *AISTATS*, 2018.
- Arbel, M., Sutherland, D. J., Binkowski, M., and Gretton, A. On gradient regularizers for MMD GANs. In *NeurIPS*, 2018.
- Arratia, R. and Gordon, L. Tutorial on large deviations for the binomial distribution. *Bulletin of Mathematical Biology*, 51:125–131, 1989.
- Barron, A. and Sheu, C.-H. Approximation of density functions by sequences of exponential families. *Annals of Statistics*, 19(3):1347–1369, 1991.
- Bellemare, M. G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., and Munos, R. The Cramer distance as a solution to biased Wasserstein gradients, 2017.
- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying MMD GANs. In *ICLR*, 2018.
- Broomhead, D. S. and Lowe, D. Radial basis functions, multi-variable functional interpolation and adaptive networks. Technical report, Royal Signals and Radar Establishment Malvern (United Kingdom), 1988.
- Canu, S. and Smola, A. J. Kernel methods and the exponential family. *Neurocomputing*, 69:714–720, 2006.
- Chwialkowski, K., Strathmann, H., and Gretton, A. A kernel test of goodness of fit. In *ICML*, 2016.
- Devroye, L. and Györfi, L. *Nonparametric Density Estimation: The  $L_1$  View*. John Wiley and Sons, 1985.
- Dheeru, D. and Karra Taniskidou, E. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using Real NVP. In *ICLR*, 2017.
- Fukumizu, K. Exponential manifold by reproducing kernel Hilbert spaces. In Gibilisco, P., Riccomagno, E., Rogantin, M.-P., and Winn, H. (eds.), *Algebraic and Geometric Methods in Statistics*, pp. 291–306. Cambridge University Press, 2009.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. Masked autoencoder for distribution estimation. In *ICML*, 2015.
- Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *JMLR*, 6(Apr):695–709, 2005.
- Janzamin, M., Sedghi, H., and Anandkumar, A. Score function features for discriminative learning: Matrix and tensor framework, 2014.
- Jean, N., Xie, S., and Ermon, S. Semi-supervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. In *NeurIPS*, 2018.
- Jimenez Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *ICML*, 2015.
- Jitkrittum, W., Xu, W., Szábo, Z., Fukumizu, K., and Gretton, A. A linear-time kernel goodness-of-fit test. In *NIPS*, 2017.
- Jitkrittum, W., Kanagawa, H., Szábo, Z., Sangkloy, P., Hays, J., Schölkopf, B., and Gretton, A. Informative features for model comparison. In *NeurIPS*, 2018.
- Kingma, D. P. and Ba, J. L. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kingma, D. P. and LeCun, Y. Regularized estimation of image statistics by score matching. In *NIPS*, 2010.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. A tutorial on energy-based learning. In *Predicting structured data*. MIT Press, 2006.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. MMD GAN: Towards deeper understanding of moment matching network. In *NIPS*, 2017.
- Liao, J. G. and Berg, A. Sharpening Jensen’s inequality. *The American Statistician*, 2018.
- Lyu, S. Interpretation and generalization of score matching. In *Uncertainty in Artificial Intelligence*, UAI ’09, 2009.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In *NIPS*, 2017.
- Saremi, S., Mehrjou, A., Schölkopf, B., and Hyvärinen, A. Deep energy estimator networks, 2018.

- Sasaki, H., Hyvärinen, A., and Sugiyama, M. Clustering via mode seeking by direct estimation of the gradient of a log-density. In *ECML/PKDD*, volume Lecture Notes in Computer Science Part III, vol. 8726, pp. 19–34, 2014.
- Shi, J. and Malik, J. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- Sriperumbudur, B., Fukumizu, K., Gretton, A., Hyvärinen, A., and Kumar, R. Density estimation in infinite dimensional exponential families. *JMLR*, 18(1):1830–1888, 2017.
- Strathmann, H., Sejdinovic, D., Livingstone, S., Szabó, Z., and Gretton, A. Gradient-free Hamiltonian Monte Carlo with efficient kernel exponential families. In *NIPS*, 2015.
- Sun, S., Kolar, M., and Xu, J. Learning structured densities via infinite dimensional exponential families. In *NIPS*, 2015.
- Sutherland, D. J., Strathmann, H., Arbel, M., and Gretton, A. Efficient and principled score estimation with Nyström kernel exponential families. In *AISTATS*, 2018.
- Uria, B., Murray, I., and Larochelle, H. Rnade: The real-valued neural autoregressive density-estimator. In *NIPS*, 2013.
- van den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *ICML*, 2016.
- Wasserman, L. *All of Nonparametric Statistics*. Springer, 2006.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. In *AISTATS*, 2016.