

# 000 001 002 003 004 005 NEURAL POSTERIOR ESTIMATION WITH LATENT BASIS 006 EXPANSIONS 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024

## ABSTRACT

025 Neural posterior estimation (NPE) is a likelihood-free amortized variational in-  
026 ference method that approximates projections of the posterior distribution. To  
027 date, NPE variational families have been either simple and interpretable (such as  
028 the Gaussian family) or highly flexible but black-box and potentially difficult to  
029 optimize (such as normalizing flows). In this work, we parameterize variational  
030 families via basis expansions of the latent variables. The log density of our varia-  
031 tional distribution is a linear combination of latent basis functions (LBFs), which  
032 may be fixed a priori or adapted to the problem class of interest. Our training  
033 and inference procedures are computationally efficient even for problems with  
034 high-dimensional latent spaces, provided only a low-dimensional projection of  
035 the posterior is of interest, owing to NPE’s automatic marginalization capabili-  
036 ties. In numerous inference problems, the proposed variational family exhibits  
037 better performance than existing variational families used with NPE, including  
038 mixtures of Gaussians (mixture density networks) and normalizing flows, as well  
039 as outperforming an existing basis expansion method for variational inference.  
040

## 041 1 INTRODUCTION

042 Neural Posterior Estimation (NPE) is an increasingly popular approach to Bayesian inference (Pa-  
043 pamakarios & Murray, 2016; Cranmer et al., 2020; Dax et al., 2021; Ward et al., 2022). In NPE,  
044 a neural network is trained exclusively with synthetic data—latent variables drawn from the prior  
045 paired with its observations—to learn the inverse mapping from an observation to its latent variables.  
046 Once trained, this network can produce posterior approximations for real data in a single forward  
047 pass. In contrast to traditional (ELBO-based) variational inference, NPE does not require likelihood  
048 evaluations. Furthermore, when the generative model contains both parameters of interest and nui-  
049 sance variables, NPE can automatically marginalize over the nuisance parameters during training:  
050 by simulating complete data and then discarding the nuisance variables to create training pairs, the  
051 method infers posterior projections for the parameters of interest (Ambrogioni et al., 2019).  
052

053 Despite these advantages, NPE shares with traditional ELBO-based variational inference a funda-  
054 mental trade-off between the flexibility of the variational family and the tractability of optimization.  
055 Simple variational families like Gaussians enable stable optimization, yet often lack the expressiveness  
056 needed for complex posterior geometries. More flexible alternatives such as Gaussian mixture  
057 models and normalizing flows add flexibility but create potentially difficult optimization landscapes  
058 with shallow local optima. Recent theoretical work (McNamara et al., 2024a) has established con-  
059 ditions for global convergence in NPE that hold for simple Gaussian variational families, but these  
060 results do not extend to the more expressive families commonly deployed in practice.  
061

062 In this work, we propose a variational family specialized for NPE, which leverages NPE’s automatic  
063 marginalization capabilities and likelihood-free nature. Unlike standard variational inference, where  
064 nuisance variables must be modeled in the variational distribution, NPE applications typically  
065 require posteriors over just a few scientifically relevant parameters: high-dimensional posterior  
066 samples do not directly aid in interpretation but must instead be post-processed to estimate low-  
067 dimensional interpretable quantities. Because NPE does not require the calculation of likelihoods,  
068 this post-processing can often be incorporated directly into the Bayesian model, resulting in a low-  
069 dimensional latent space of interest. In these low-dimensional settings, even numerical integration is  
070 computationally feasible, freeing us from the usual requirement of closed-form normalization.  
071

We leverage this freedom by parameterizing the log density of variational distributions through latent basis expansions: a neural network processes observations to produce coefficients for linear combinations of basis functions over the latent space. This approach yields distributions in the exponential family—among the most flexible classes available (Pacchiardi & Dutta, 2022; Khemakhem et al., 2020; Sriperumbudur et al., 2017)—while maintaining favorable optimization properties. The resulting method, which we refer to as Latent Basis Function NPE (LBF-NPE), optimizes over the class of all exponential families of a fixed dimension  $K$  by adaptively fitting the basis functions, denoted  $s_\psi(z) \in \mathbb{R}^K$ . Simultaneously, amortization is performed by fitting a separate network  $f_\phi(x) \in \mathbb{R}^K$  that maps observations  $x$  to coefficients of these basis functions (Section 3).

In Section 4, we introduce and analyze several variants of LBF-NPE. For a variant with fixed basis functions, such as B-splines or wavelets, optimization is convex despite the log-normalizer, providing stable training that has proven elusive for more complex variational families, and ensuring stable convergence to global optima under the conditions presented by McNamara et al. (2024a). Alternatively, both coefficients  $f_\phi(x)$  and basis functions  $s_\psi(z)$  can be fitted jointly through alternating optimization that exploits marginal convexity in each component. We employ stereographic projection reparameterization to address identifiability issues in this adaptive setting, constraining outputs to the unit hypersphere and stabilizing training.

We demonstrate superior performance of LBF-NPE across diverse inference problems, from synthetic benchmarks to real scientific applications (Section 6). LBF-NPE consistently converges to global optima on multimodal problems where mixture density networks converge to shallow local minima. LBF-NPE with just 20 basis functions achieves order-of-magnitude improvements in KL divergence over both MDNs and normalizing flows on complex 2D posteriors. The method successfully captures multimodal posteriors in astronomical object detection and substantially outperforms MDN baselines on cosmological redshift estimation using the LSST DESC DC2 survey dataset.

## 2 BACKGROUND

### 2.1 ELBO-BASED VARIATIONAL INFERENCE

In variational inference (VI), numerical optimization is used to select an approximation  $q(z)$  of the posterior distribution of some model  $p(z, x)$  on observables  $x$  and latent variables  $z$ . The most common variational objective, the evidence lower bound (ELBO), targets minimization of the reverse KL divergence (Blei et al., 2017; Zhang et al., 2019; Kingma & Welling, 2019) by constructing the variational quantity

$$\text{ELBO}(\eta) := \mathbb{E}_{q(z; \eta)} \log \left( \frac{p(z, x)}{q(z; \eta)} \right) \leq \log p(x) \quad (1)$$

and performing maximization in  $\eta$  for a fixed choice of  $x$ . Optimization of the ELBO continues to pose a longstanding difficulty in VI. Firstly, even simple variational families such as the Gaussian can exhibit problematic optimization landscapes without a careful choice of parameterization. Targeting the ELBO, even in the non-amortized setting, is generally a nonconvex problem (Liu et al., 2023a; Domke, 2020; Domke et al., 2023). Secondly, and most significantly in our setting, ELBO-based variational inference does not provide a way to marginalize over nuisance latent variables. For a model  $p(z, \xi, x)$  on latent variables  $\{z, \xi\}$  and observed variables  $x$ , ELBO-based methods typically cannot compute the quantity  $p(z, x) = \int p(z, \xi, x) d\xi$  required to target their objective function.

### 2.2 NEURAL POSTERIOR ESTIMATION

In amortized variational inference, the shared parameters of a deep neural network define variational approximations for arbitrary observations  $x$  (Ganguly et al., 2024; McNamara et al., 2024a). Precisely, a variational approximation for latent variable  $z$  is given by  $q(z; \eta)$  with parameters  $\eta$ . Rather than fitting these parameters separately for each  $x$ , the amortized approach defines  $\eta = f_\phi(x)$  for a neural network  $f_\phi$  with parameters  $\phi$ . This inference network, once fit, yields the variational posterior  $q(z; f_\phi(x))$  for arbitrary  $x$  by a single forward pass (Kingma & Welling, 2019; Ambrogioni et al., 2019). Neural posterior estimation (NPE) (Papamakarios & Murray, 2016) targets an expectation of the forward KL divergence for amortized VI:

$$\tilde{\mathcal{L}}_{\text{NPE}}(\phi) = \mathbb{E}_{p(x)} \text{KL}(p(z | x) || q(z | x)). \quad (2)$$

108 Here, the integral over  $p(x)$ , the marginal of the model  $p(z, \xi, x)$ , indicates that the NPE objective  
 109 averages over all possible draws from the generative process instead of averaging over observations  $x$   
 110 from some finite training set. The objective is equivalent (up to a constant) to

$$\mathcal{L}_{\text{NPE}}(\phi) = -\mathbb{E}_{p(z,x)} \log q(z; f_\phi(x)), \quad (3)$$

113 where  $q_\phi(z; \eta) \equiv q(z; f_\phi(x))$  for any  $x$ . Equation (3) admits unbiased estimation of its gradient with  
 114 stochastic draws  $(z, x) \sim p(z, \xi, x)$  from the joint model, readily obtained by ancestral sampling, even  
 115 in the presence of nuisance latent variables  $\xi$ . Explicitly, we can sample  $(z, x) \sim p(z, x)$  by simulating  
 116 an entire sequence  $p(z)p(\xi_1 | z)p(\xi_2 | \xi_1) \cdots p(\xi_L | \xi_{L-1})p(x | \xi_L)$  and discarding variables that  
 117 are not the target for inference. The expected forward KL objective has been independently derived  
 118 and analyzed in several related works (Bornschein & Bengio, 2015; Ambrogioni et al., 2019).

### 3 LBF-NPE: BASIS EXPANSIONS FOR AMORTIZED LOG DENSITY ESTIMATION

123 We propose an amortized method to fit complex multimodal variational distributions, called Latent  
 124 Basis Function NPE (LBF-NPE). The method fits a basis function network  $s_\psi : z \mapsto \mathbb{R}^K$ , which  
 125 evaluates  $K$  basis functions  $[s_\psi^{(1)}, \dots, s_\psi^{(K)}]$  for any point  $z$  in the latent space. The inference  
 126 network  $f_\phi : x \mapsto \eta \in \mathbb{R}^K$  maps observations to coefficients of the basis functions. The number  
 127 of basis functions  $K$  is fixed ahead of time, and larger  $K$  may be used to increase expressivity (see  
 128 Appendix E). The variational parameters to be fit are neural network weights  $\phi$  and  $\psi$ , for each of the  
 129 networks  $f_\phi$  and  $s_\psi$ . Below, we first construct an exponential family defined by the basis functions  
 130 (Section 3.1) and then give the fitting routine for  $f_\phi$  and  $s_\psi$  (Section 3.2).

#### 3.1 THE VARIATIONAL FAMILY

132 Fix  $x \in \mathcal{X}$ . Let  $I \subseteq \mathbb{R}^d$  denote the latent space and let  $\{s_\psi^{(i)}(z)\}_{i=1}^K$  be a collection of basis functions  
 133 defined by parameters  $\psi$ . Selecting more basis functions (larger  $K$ ) leads to a more expressive  
 134 variational distribution, but also a higher-dimensional optimization problem. For a value  $z \in I$ , let  
 135  $s_\psi(z) = [s_\psi^{(1)}, \dots, s_\psi^{(K)}]^\top \in \mathbb{R}^K$ . Our variational family is parameterized by coefficients  $\eta \in \mathbb{R}^K$ ,  
 136 and has the density function

$$q(z; \eta) \propto h(z) \exp(\eta^\top s_\psi(z)). \quad (4)$$

137 We aim to select  $\eta$  and  $\psi$  such that  $q(z; \eta) \approx p(z | x)$ . The log density is

$$\log q(z; \eta) = \log h(z) + \eta^\top s_\psi(z) - C, \quad (5)$$

138 where  $C$  is the log of the normalizing constant. This variational family (Equation 4) is an exponential  
 139 family (Wainwright & Jordan, 2008; Srivastava et al., 2014). The vector of basis functions  $s_\psi(z)$  is  
 140 the sufficient statistic,  $\eta$  is the natural parameter vector, and  $h(z)$  is any finite base measure on the  
 141 latent space  $I$  (i.e.,  $\int_I h(z) dz$  is finite).

142 We represent the  $K$ -dimensional quantity  $s_\psi(z)$  as the output of a deep neural network with input  
 143  $z$ . Because the number and form of the basis functions  $s_\psi(z)$  are arbitrary, the expressivity of this  
 144 variational family is far greater than that of “classical” exponential families, such as the Gaussian  
 145 family. As  $K \rightarrow \infty$ , the set of *all* exponential family distributions is arbitrarily rich: any distribution  
 146 can be represented as an infinite-dimensional exponential family distribution (Khemakhem et al.,  
 147 2020; Sriperumbudur et al., 2017).

#### 3.2 THE AMORTIZED VARIATIONAL OBJECTIVE & GRADIENT ESTIMATOR

148 The formulation of Section 3.1 is *non-amortized*: it requires selecting a single  $\eta$  for  $q(z; \eta)$  to  
 149 approximate the posterior for a single  $x$ . In NPE, we consider the amortized problem, where we  
 150 define the posterior for arbitrary  $x$ . We set  $\eta = f_\phi(x)$ , and thus require fitting two separate networks,  
 151  $f_\phi$  and  $s_\psi$ . We fit the variational parameters,  $\phi$  and  $\psi$ , by minimizing the NPE objective function,  
 152 given in general in Section 2.2. Our specific variational objective (up to constants) is thus

$$L_{\text{LBF-NPE}}(\phi, \psi) = -\mathbb{E}_{p(z,x)} \left( f_\phi(x)^\top s_\psi(z) - \log \left( \int \exp(f_\phi(x)^\top s_\psi(\tilde{z})) h(\tilde{z}) d\tilde{z} \right) \right), \quad (6)$$

162 where the log-normalizer takes the form of the log of an integral with respect to the base measure  
 163  $h$ . Estimation of the log-normalizer cannot be done unbiasedly by Monte Carlo sampling due to  
 164 the Jensen gap (Adil Khan et al., 2015). For training, we only require stochastic gradients, which  
 165 can be computed using importance sampling. Focusing on the log-normalizer for now, we let  
 166  $k_{\phi,\psi}(x, z) = f_\phi(x)^\top s_\psi(z)$ . We suppress the dependence on  $x$  for now, as it is fixed in the integral.  
 167 Let  $J(\phi, \psi) := \int \exp k_{\phi,\psi}(\tilde{z}) dh(\tilde{z})$ . The gradient can be computed by estimating an expectation  
 168 with respect to an exponentially tilted transformation of  $h$ . Let  $\nabla_{\phi,\psi}$  denote the gradient with respect  
 169 to either  $\phi$  or  $\psi$ . Then, we have

$$\nabla_{\phi,\psi} \log J(\phi, \psi) = \frac{1}{J(\phi, \psi)} \cdot \nabla_{\phi,\psi} J(\phi, \psi) \quad (7)$$

$$= \frac{1}{J(\phi, \psi)} \cdot \int [\nabla_{\phi,\psi} k_{\phi,\psi}(\tilde{z})] \cdot \exp k_{\phi,\psi}(\tilde{z}) h(\tilde{z}) d\tilde{z} \quad (8)$$

$$= \int [\nabla_{\phi,\psi} k_{\phi,\psi}(\tilde{z})] \cdot \left( \frac{\exp k_{\phi,\psi}(\tilde{z}) h(\tilde{z})}{\int \exp k_{\phi,\psi}(z') h(z') dz'} \right) d\tilde{z} \quad (9)$$

$$=: \int [\nabla_{\phi,\psi} k_{\phi,\psi}(\tilde{z})] \cdot q_{\phi,\psi}(\tilde{z}) d\tilde{z}, \quad (10)$$

180 where we recognize an expectation with respect to  $q_{\phi,\psi}$ , an exponentially tilted density that depends  
 181 on the current values of  $\phi, \psi$ . This integral can be estimated by the use of self-normalized importance  
 182 sampling (SNIS) with a proposal distribution  $r(\tilde{z})$  (Owen, 2013):

$$\int [\nabla_{\phi,\psi} k_{\phi,\psi}(\tilde{z})] q_{\phi,\psi}(\tilde{z}) d\tilde{z} \approx \sum_{j=1}^P \frac{[\nabla_{\phi,\psi} k_{\phi,\psi}(\tilde{z}_i)] w(\tilde{z}_i)}{\sum_{k=1}^P w(\tilde{z}_k)},$$

188 where  $w(z) = \frac{\exp(k_{\phi,\psi}(z)h(z))}{r(z)}$  and  $\tilde{z}_1, \dots, \tilde{z}_K \stackrel{iid}{\sim} r$ . The gradient estimator is thus biased for the true  
 189 gradient, similar to other gradient estimators targeting the forward KL from the family of “wake-sleep”  
 190 algorithms, but consistent as  $P \rightarrow \infty$  (Le et al., 2019; Bornschein & Bengio, 2015; McNamara et al.,  
 191 2024b). Algorithm 1 details our gradient computation procedure.

---

**Algorithm 1:** Gradient Computation for LBF-NPE
 

---

195 **Inputs:** Sampling model  $p(z, x)$ ; networks  $f_\phi$  and  $s_\psi$ ; proposal distribution  $r(z)$ .  
 196 Sample batch  $\{(z_i, x_i)\}_{i=1}^B \stackrel{iid}{\sim} p(z, x)$   
 197 /\* Gradient for log-normalizer \*/  
 198 Propose  $\tilde{z}_1, \dots, \tilde{z}_P \sim r$   
 199 Compute  $k_j^i := k_{\phi,\psi}(x_i, \tilde{z}_j) = f_\phi(x_i)^\top s_\psi(\tilde{z}_j)$ ,  $i \in [B], j \in [P]$   
 200 Compute unnormalized weights  $w_j^i = \exp(k_j^i h(\tilde{z}_j))$ ,  $i \in [B], j \in [P]$   
 201 Compute  $U_{\phi,\psi}^i = \sum_{j=1}^P \frac{w_j^i \nabla_{\phi,\psi} k_j^i}{\sum_{j'} w_{j'}^i}$ ,  $i \in [B]$   
 202 /\* Gradient non-tilted inner product \*/  
 203 Compute  $V_{\phi,\psi}^i = \nabla_{\phi,\psi} k_{\phi,\psi}(x_i, z_i) = f_\phi(x_i)^\top s_\psi(z_i)$   
 204 /\* Compute combined gradient \*/  
 205 Return  $\hat{\nabla}_{\phi,\psi} = -\frac{1}{n} \sum_{i=1}^n V_{\phi,\psi}^i - U_{\phi,\psi}^i$ 


---

**4 VARIANTS & PROPERTIES OF LBF-NPE**

213 We now motivate the construction of the LBF-NPE variational family by examining aspects of its  
 214 optimization routine. As a result of parameterizing and targeting the log-density (cf. 5, 7), both  
 215 our construction and optimization routine depend entirely on the inner product  $f_\phi(x)^\top s_\psi(z)$ . We  
 elaborate on some key properties and variants of LBF-NPE that stem from this observation.

216 4.1 AFFINE GRADIENTS  
217218 In Section 3.2, we showed that the gradient of the objective for LBF-NPE depends only on the  
219 gradient of  $k_{\psi, \phi}(x, z) = f_{\phi}(x)^{\top} s_{\psi}(z)$ , via the relation

220 
$$\nabla_{\phi, \psi} L_{\text{LBF-NPE}}(\phi, \psi) = -\mathbb{E}_{p(z, x)} [\nabla_{\phi, \psi} k_{\psi, \phi}(x, z) - \mathbb{E}_{q_{\phi, \psi}(\tilde{z})} [\nabla_{\phi, \psi} k_{\phi, \psi}(x, \tilde{z})]] \quad (11)$$
  
221

222 where  $q_{\phi, \psi}$  denotes the exponentially tilted density constructed in Section 3.2. Accordingly, the form  
223 of gradient updates for the LBF-NPE procedure is extremely simple; in fact, holding  $\psi$  constant,  
224 the gradient with respect to  $\phi$  is that of an affine function of the network outputs  $f_{\phi}$ . The same  
225 relationship holds when taking gradients for  $\psi$  holding  $\phi$  constant. Targeting such simple functions for  
226 optimization, besides being simple to implement, benefits from a convex formulation (see Section 4.2  
227 below). The invariance of the inner product  $k_{\psi, \phi}$  under arbitrary rescalings of  $f$  and  $s$ , on the other  
228 hand, complicates optimization: this motivates a variant of LBF-NPE that reparameterizes outputs to  
229 unit norm (see Section 4.4).  
230

## 231 4.2 CONVEXITY

232 McNamara et al. (2024a) show that neural posterior estimation (NPE) optimizes a convex functional  
233 objective function provided that the variational family is log-concave in  $f$ , the inference network.  
234 This ensures the forward KL objective of NPE (cf. Equation 3) is convex in  $f$ . Recent advances in  
235 the study of wide networks via the neural tangent kernel (NTK) (Jacot et al., 2018) have shown that  
236 fitting network parameters to minimize convex loss functionals (e.g., mean squared error) follows  
237 kernel gradient descent dynamics to a global optimum in the infinite-width limit (Jacot et al., 2018;  
238 McNamara et al., 2024a).239 The amortized forward KL objective function that we target (cf. Equation 6) benefits from these same  
240 properties. For an arbitrary collection of basis functions, the objective function  $L_{\text{LBF-NPE}}$  remains a  
241 convex functional in  $f$ . Likewise, for fixed  $f$ ,  $L_{\text{LBF-NPE}}$  is a convex functional in  $s$ . We formalize this  
242 in the proposition below.243 **Proposition 1.** *The functional*

244 
$$L(f, s) = -\mathbb{E}_{p(z, x)} \left( f(x)^{\top} s(z) - \log \left( \int \exp (f(x)^{\top} s(\tilde{z})) dh(\tilde{z}) \right) \right)$$
  
245

246 is marginally convex in the argument  $f$  and  $s$ , respectively.247 A proof and additional discussion are provided in Appendix B. Proposition 1 shows that in the case  
248 where either  $f$  or  $s$  are fixed, the resulting functional that is optimized is fully convex, rather than just  
249 marginally so. This observation motivates a variant of LBF-NPE where the basis functions are fixed  
250 a priori (see Section 4.3).  
251252 4.3 FIXED BASIS FUNCTIONS  
253254 Rather than adaptively fitting basis functions  $s_{\psi}$ , the practitioner may simply use a fixed basis defined  
255 ahead of time. This approach is motivated by the convexity of the resulting functional in  $f$ , as  
256 well as the approaches of related work based on basis expansion parameterizations, which use fixed  
257 orthonormal eigenfunctions (cf. Section 5). In this variant of LBF-NPE, the objective function  
258  $L(\phi, \psi)$  collapses to the marginal  $L(\phi)$  for optimization. As we elaborate in Section 4.2, LBF-NPE  
259 has a convex formulation in this setting, which empirically results in advantageous optimization  
260 trajectories relative to competing methods (we demonstrate this in Section 6.1).  
261262 Several choices of basis may be of interest to practitioners. EigenVI, a related basis-expansion method  
263 for VI (cf. Section 5), utilizes a (truncated) orthonormal basis of eigenfunctions, such as Bernstein,  
264 Legendre, or Hermite polynomials (Cai et al., 2024). Selecting a large  $K$  improves faithfulness to the  
265 complete basis, but doing so increases the dimension of the optimization problem, exponentially so  
266 in multiple dimensions. Further, as generally such basis functions are *global* (i.e., nonzero on all of  
267 the latent space  $I$ ), in this design *every* basis function contributes to the density value  $q(z)$  at every  
268 point  $z$ ; this may make it difficult to control the local behavior of the fitted density.269 An alternative approach is to model  $\log q(z; \eta) = \log q(z; f_{\phi}(x))$  via a *local* basis expansion. We  
270 specialize to B-splines (Appendix A.1) and wavelets (Appendix A.2) in our experiments, two rich

270 families that we recommend for practitioners. In this framework, each basis function is nonzero only  
 271 in a small neighborhood of the latent support. Locality of the basis functions simplifies optimization  
 272 by inducing a sparser problem than a set of *global* basis functions would. For a single Monte Carlo  
 273 draw  $(z^*, x^*) \sim p(z, x)$ , the gradient  $-\nabla_\eta \log q(z^*; \eta) |_{\eta=f_\phi(x^*)}$  is nonzero at only a few indices  
 274 because many basis functions are zero at any given  $z^*$ .

#### 276 4.4 REDUCING DEGENERACY THROUGH STEREOGRAPHIC PROJECTION

277 As noted in Section 4.1, both gradients and the log-density itself only depend on the inner product  
 278  $f_\phi(x)^\top s_\psi(z)$ . Adaptively learning both the inference network  $f_\phi$  and basis function network  $s_\psi$   
 279 thus suffers from an inherent lack of identifiability: different rescalings or rotations of the vectors  
 280 defined by  $f_\phi$  and  $s_\psi$  can lead to the same loss function values, since the loss function (Equation 6)  
 281 only depends on the inner product. To mitigate this degeneracy, we propose a variant of LBF-NPE  
 282 that uses stereographic projection reparameterization to normalize the output tensor onto the unit  
 283 hypersphere. Specifically, for a  $K$ -dimensional coefficient or basis function vector, we construct  
 284 a neural network that outputs a  $(K - 1)$ -dimensional vector  $u$ . We then apply the stereographic  
 285 projection reparameterization:  $y = \left( \frac{2u}{1+\|u\|^2}, \frac{1-\|u\|^2}{1+\|u\|^2} \right)$  which maps  $u \in \mathbb{R}^{K-1}$  to a unit vector  $y \in$   
 286  $\mathbb{R}^K$  such that  $\|y\| = 1$ . This normalization mitigates identifiability issues, and the reparameterization  
 287 yields strong results in our experiments (see Appendix D for additional discussion). Our loss function  
 288 takes the following form when we apply this reparameterization:  
 289

$$290 \hat{L}_{\text{LBF-NPE}}(\phi, \psi) = \mathbb{E}_{p(z, x)} \left( -w \hat{f}_\phi(x)^\top \hat{s}_\psi(z) + \log \left( \int \exp \left( w \hat{f}_\phi(x)^\top \hat{s}_\psi(\tilde{z}) \right) d\tilde{z} \right) \right), \quad (12)$$

294 where  $\hat{f}_\phi(\cdot)$  and  $\hat{s}_\psi(\cdot)$  are reparameterized network outputs, and  $w$  is a fixed scaling factor.  
 295

## 296 5 RELATED WORK

297 Exponential family distributions are a common class of distributions for both traditional variational  
 298 inference and NPE. In the simplest cases, Gaussian, Bernoulli, and other “simple” exponential  
 299 families are used (Liu et al., 2023a; Cranmer et al., 2020; Blei et al., 2017). Typically, however, these  
 300 distributions are not parameterized in canonical form (where  $\eta = f_\phi(x)$  is the natural parameter of the  
 301 family). However, for NPE, McNamara et al. (2024a) recommends using the canonical parameterization  
 302 even for simple families such as the Gaussian to benefit from convex loss (Section 4.2). General  
 303 exponential families parameterized by neural networks were first proposed in Pacchiardi & Dutta  
 304 (2022) to represent the *likelihood* function for likelihood-free settings. Akin to Approximate Bayesian  
 305 Computation (ABC) methods, this approach aims to learn low-dimensional summary statistics of  $x$   
 306 to represent the likelihood, and subsequently performs inference with potentially expensive MCMC  
 307 or ABC routines. We compare to this approach in Appendix E.6. To our knowledge, LBF-NPE is  
 308 the first method to utilize neural exponential families to represent the posterior distribution and to  
 309 use this family within the amortized inference setting. LBF-NPE is also unique in exploiting the low  
 310 dimensionality of the posterior projections of interest.  
 311

312 Parameterizing variational distributions via basis expansions is a relatively new line of research; a  
 313 recent non-amortized approach, EigenVI (Cai et al., 2024), presents an algorithm for optimizing a  
 314 score-based divergence with a variational family parameterized via a linear combination of orthogonal  
 315 eigenfunctions. Key limitations of this approach are i) the lack of amortization and ii) the necessity of  
 316 utilizing orthogonal, fixed eigenfunctions as the basis: truncation of these bases necessarily introduces  
 317 approximation error. In our case,  $s_\psi$  is unrestricted: the basis functions can be arbitrary, and so a  
 318 fixed number  $K$  may be sufficient for some classes of posteriors (cf. Section 6.2).

319 Mixtures of Gaussians and normalizing flows are other common choices of variational families for  
 320 NPE (Gershman et al., 2012; Papamakarios & Murray, 2016; Papamakarios et al., 2021; Rezende  
 321 & Mohamed, 2015). Although more flexible than simple exponential families, these parameteriza-  
 322 tions may suffer from convergence to shallow local optima during optimization (cf. Section 6.1).  
 323 We compare LBF-NPE to mixtures of Gaussians, normalizing flows, as well as EigenVI in our  
 experiments.

324 

## 6 EXPERIMENTS

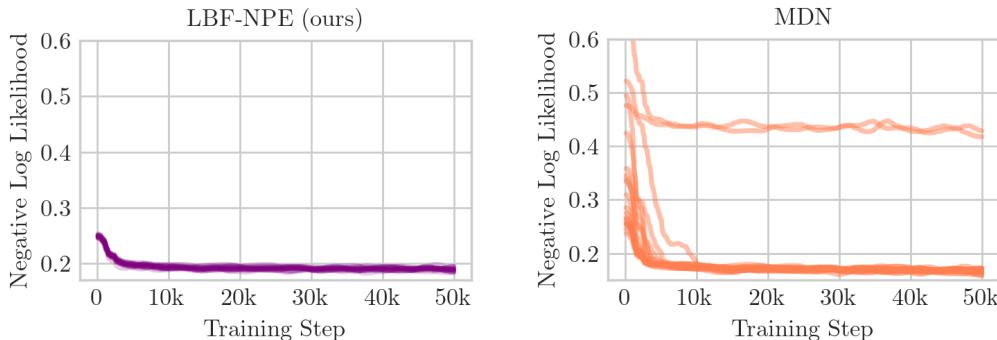
326 In numerical experiments, we fit a variety of complex posterior distributions using LBF-NPE. In  
 327 Sections 6.1 and 6.4, we infer one-dimensional posterior projections using the variant of our method  
 328 with fixed basis functions, whereas in Sections 6.2 and 6.3, we infer two-dimensional posterior  
 329 projections using adaptive basis functions. Additional details about each of these experiments appear  
 330 in Appendix D.

331 NPE with various alternative variational families serves as our primary benchmarks; we can compare  
 332 to these methods quantitatively by assessing the NPE objective with each choice of variational  
 333 distribution. We benchmark NPE with variational families based on mixture density networks  
 334 (MDNs), RealNVP, and neural spline flows (Papamakarios & Murray, 2016; Durkan et al., 2019). In  
 335 addition to the results in this section, additional results appear in Appendix E, including results from  
 336 comparisons to two non-NPE-based variational inference methods: EigenVI (Appendix E.5) and a  
 337 score-matching neural-likelihood-based method for likelihood-free inference (Appendix E.6).

339 

### 6.1 TOY EXAMPLE: SINUSOIDAL LIKELIHOOD

340 We first exhibit the advantages of LBF-NPE’s convex variational objective by demonstrating consistent  
 341 convergence on a highly multimodal problem with fixed basis functions (cf. Section 4.3). The model  
 342 draws an angle  $z \sim \text{Unif}[0, 2\pi]$  followed by  $x \mid z \sim \mathcal{N}(\sin(2z), \sigma^2)$  for fixed  $\sigma^2 = 1$ . The  
 343 exact posteriors  $p(z \mid x)$  have up to four modes, depending on the realization  $x$ . We compare  
 344 LBF-NPE, using a fixed collection of 14 B-spline functions of degree two on a mesh of  $[0, 2\pi]$   
 345 (see Appendix A.1 for additional detail on B-splines), and a mixture density network (MDN)  
 346 (Bishop, 1994; Papamakarios & Murray, 2016), using a mixture of five Gaussian distributions. These  
 347 variational distributions both have 14 distributional parameters for each observation  $x$ . Additional  
 348 experimental details are given in Appendix D. Figure 1 shows that for 20 different runs of the  
 349 optimization routine, LBF-NPE consistently converges to the same solution, whereas the MDN  
 350 sometimes converges to a suboptimal local optimum. Visualizations of posterior approximations  
 351 from both methods are provided in Appendix E.



364 Figure 1: Negative log likelihood of our method and the MDN. Each model is trained with 20  
 365 different random seeds, and the records are smoothed using a Gaussian filter with  $\sigma = 20.0$ .  
 366

368 

### 6.2 COMPLEX MULTIVARIATE REPRESENTATIONS IN 2D

370 We showcase LBF-NPE on three test problems in two dimensions, named “banded”, “ring”, and  
 371 “spiral,” and visualized in the left column of Figure 2. Each model consists of two-dimensional latent  
 372 variable  $z \in \mathbb{R}^2$  and observation  $x \in \mathbb{R}$ , and in some cases nuisance latent variables as well. Further  
 373 details of the generative processes for these three problems are provided in Appendix D (Sections D.2  
 374 to D.4).

375 LBF-NPE is able to approximate these complex posteriors nearly perfectly using only 20 adaptive  
 376 basis functions  $s_\psi$ . Both the amortization function  $f_\phi$  and the basis functions  $s_\psi$  are parameterized  
 377 using deep neural networks. For additional implementation details, we refer the reader to Appendix D.  
 We follow Algorithm 1 and evaluate both the variational posterior  $q(z; f_\phi(x), s_\psi(z))$  and the exact

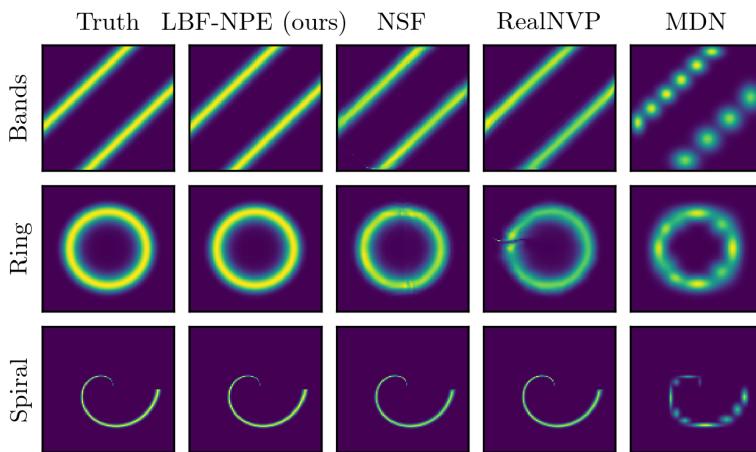


Figure 2: Example posteriors of three problems in two dimensions. NSF refers to Neural Spline Flow.

Table 1: Forward/reverse KL divergence and NLL of LBF-NPE (ours), NSF (Neural Spline Flow), RealNVP, and MDN on three 2D test cases. Lower values indicate better posterior approximation.

Forward KL Divergence				Reverse KL Divergence				NLL				
	LBF-NPE	NSF	RealNVP	MDN	LBF-NPE	NSF	RealNVP	MDN	LBF-NPE	NSF	RealNVP	MDN
Bands	<b>0.0048</b>	0.016	0.015	0.182	<b>0.0014</b>	0.0099	0.011	0.156	<b>-0.060</b>	0.151	0.157	1.389
Ring	<b>0.0054</b>	0.017	0.024	0.205	<b>0.0027</b>	0.013	0.014	0.204	<b>0.030</b>	0.621	0.733	1.031
Spiral	<b>0.187</b>	0.201	0.545	0.948	<b>0.188</b>	0.322	0.666	1.973	0.838	<b>0.727</b>	0.859	2.788

posterior  $p(z | x)$  on a fine mesh grid that covers the support of the posterior. Qualitative results appear in Figure 2 and quantitative results appear in Table 1 for held-out test points. LBF-NPE outperforms the MDN and multiple types of normalizing flow on both metrics. We provide additional visualizations of the variational approximations found by LBF-NPE and its competitors in Appendix E.

### 6.3 OBJECT DETECTION

We apply LBF-NPE with adaptive basis functions to the problem of object detection in astronomical images. We use a generative model resembling the scientific model of Liu et al. (2023b). In brief, this generative process first independently samples star locations  $l_1, l_2 \sim \text{Unif}([0, 16] \times [0, 16])$  and star fluxes  $f_1, f_2 \sim \mathcal{N}(\mu, \sigma^2)$  for two objects. Afterward, a latent noise-free image  $I$  is rendered by convolving these point sources with a point-spread function (PSF). Finally, given  $I$ , the intensity of each pixel  $(j, k)$ , for  $j, k \in \{1, \dots, 16\}$  is independently drawn as  $x_{jk} \sim \text{Poisson}(I_{jk})$ , reflecting Poisson shot noise.

Figure 3 shows two examples of the noisy observations  $x$ , along with the posterior approximations for the locations of each. The posterior distribution for this problem is multimodal with a high degree of separation between modes. LBF-NPE parameterizes this shape effectively. For LBF-NPE, we further vary the number of basis functions  $K = 9, 20, 36, 64$ . Examining the form of the fitted basis functions across varying  $K$  (Appendix E.3) is particularly illustrative of the advantages of the adaptive approach to fitting the basis functions.

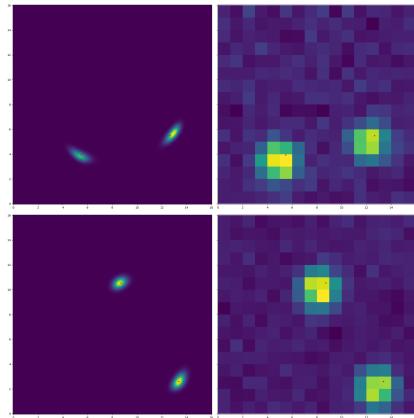


Figure 3: Two example posteriors (left) conditional on the observed images  $x$  (right). In each case, LBF-NPE correctly recovers the locations of the two objects.

432 6.4 CASE STUDY: REDSHIFT ESTIMATION  
433

434 The redshift of galaxies is a key quantity of interest as it characterizes their distances from Earth. Redshift measures the extent to  
435 which electromagnetic waves are “stretched” to redder wavelengths  
436 as objects move away from Earth. The distribution of redshifts  
437 across many objects is a powerful probe of cosmology (Malz &  
438 Hogg, 2022). Redshift estimation from photometric data (images)  
439 is referred to as photo- $z$  estimation. We extend the methodology of  
440 the Bayesian Light Source Separator (BLISS) package (Liu et al.,  
441 2023b), a state-of-the-art package for probabilistic object detection  
442 in astronomical images, for this task, by adding a redshift density  
443 estimation “head” to the existing BLISS network. To each detected object, we associate a redshift density  
444 probability density function fitted by LBF-NPE with a fixed B-spline basis family. LSST DESC DC2  
445 Simulated Sky Survey dataset (LSST DESC et al., 2021) serves as the generative model, providing  
446 simulated  $(z, x)$  pairs, where  $z$  denotes redshift and  $x$  are the astronomical images. This highly  
447 realistic dataset consists of mock catalogs of astronomical images produced by directly modeling  
448 known physical quantities of the universe using empirical priors and physics-informed modeling  
449 choices. Appendix D provides further details of the experimental setting.

450 We compare LBF-NPE to variational families based on neural spline flow (NSF) and a mixture  
451 density network (MDN), all embedded with the BLISS framework. The MDN uses five Gaussian  
452 components, in keeping with state-of-the-art work on photo- $z$  estimation (Merz et al., 2025); more  
453 than five components did not improve the quality of fit. The only difference between the two  
454 approaches is the parameterization choice of the variational family. We compute the negative log-  
455 likelihood (NLL) of a held-out test set of 153,000 astronomical objects. Table 2 shows that LBF-NPE  
456 with the B-spline variational family parameterization outperforms both the MDN and the NSF.

## 477 7 DISCUSSION &amp; LIMITATIONS

460 LBF-NPE models the log density of the variational distribution as a linear combination of expressive  
461 basis functions, which is beneficial for several reasons. First, log-space modeling results in a  
462 *multiplicative* influence of different basis functions. Regions of latent space can effectively be  
463 “zeroed-out” more easily in this context compared to performing the modeling in density space  
464 directly. Second, our model of the log density results in an unconstrained optimization problem in  
465  $f_\phi$  and  $s_\psi$ : the coefficients and basis functions may be either positive or negative, whereas other  
466 density estimation methods may require nonnegativity or other constraints on the coefficients or basis  
467 functions to obtain a valid density that integrates to unity (Cai et al., 2024; Koo & Kim, 1996; Kirkby  
468 et al., 2023).

469 Using basis expansions to parameterize variational distributions is a recent and exciting innovation  
470 in variational inference. Relative to EigenVI, LBF-NPE performs better with fewer basis functions  
471 (Appendix E.5). This may be due to removing the orthogonality constraints and adaptively fitting basis  
472 functions: by averaging across posteriors for arbitrary  $x$ , our approach ensures the basis functions are  
473 implicitly regularized, preventing overfitting to any single instance. Appendix E contains additional  
474 discussion and visualization of fitted basis functions  $s_\psi$ .

475 The main limitation of LBF-NPE is the difficulty of sampling from the variational distribution.  
476 LBF-NPE directly fits the log density of the variational distribution, but samples from this density are  
477 typically needed for inference. The low dimensionality of many NPE targets ensures that sampling  
478 is straightforward: inverse transform sampling is readily applicable (cf. Appendix C). For higher-  
479 dimensional targets, importance sampling may be adequate to estimate functionals with respect to the  
480 variational distribution.

481 Despite this sampling challenge, our approach demonstrates that basis expansion methods offer a  
482 compelling middle ground between optimization simplicity and expressivity for NPE. Future work  
483 could explore new bilevel optimization approaches (Xiao & Chen, 2025) to jointly learn adaptive basis  
484 functions alongside their coefficients. Additionally, extending our approach beyond low-dimensional  
485 targets to high-dimensional targets with simplifying structure—such as known or assumed conditional  
independencies—could broaden the applicability of basis expansion methods.

Table 2: Held-out NLL of the true redshift  $z$ .

Method	NLL
LBF-NPE	<b>-57,220</b>
NSF	-55,389
MDN	-50,648

486 REFERENCES  
487

488 M. Adil Khan, G. Ali Khan, T. Ali, and A. Kılıcman. On the refinement of Jensen’s inequality.  
489 *Applied Mathematics and Computation*, 262:128–135, 2015.

490 Luca Ambrogioni, Umut Güçlü, Julia Berezutskaya, Eva van den Borne, Yağmur Güçlütürk, Max  
491 Hinne, Eric Maris, and Marcel van Gerven. Forward amortized inference for likelihood-free  
492 variational marginalization. In *International Conference on Artificial Intelligence and Statistics*,  
493 2019.

494 Jason Ansel, Edward Yang, Horace He, et al. PyTorch 2: Faster Machine Learning Through Dynamic  
495 Python Bytecode Transformation and Graph Compilation. In *29th ACM International Conference  
496 on Architectural Support for Programming Languages and Operating Systems*. ACM, 4 2024.

497 Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine  
498 Learning Research*, 18(19):1–53, 2017.

499 500 Joshua Bengio, Nicolas Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex  
501 neural networks. In *Advances in Neural Information Processing Systems*, 2005.

502 503 Christopher M. Bishop. Mixture density networks. Technical Report NCRG/94/004, Aston University,  
504 Birmingham, UK, 1994.

505 506 David M. Blei, Alp Kucukbir, and Jon D. McAuliffe. Variational inference: A review for statisticians.  
507 *Journal of the American Statistical Association*, 112(518):859–877, 2017.

508 509 Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. In *International Conference on  
510 Learning Representations*, 2015.

511 512 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin,  
513 George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang.  
JAX: composable transformations of Python+NumPy programs, 2018.

514 515 Diana Cai, Chirag Modi, Charles C. Margossian, Robert M. Gower, David M. Blei, and Lawrence K.  
516 Saul. EigenVI: score-based variational inference with orthogonal function expansions. In *Advances  
517 in Neural Information Processing Systems*, 2024.

518 519 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for  
520 contrastive learning of visual representations. In *Proceedings of the 37th International Conference  
on Machine Learning*, 2020.

521 522 LSST Dark Energy Science Collaboration, Bela Abolfathi, Robert Armstrong, et al. DESC DC2 data  
523 release note, 2022.

524 525 Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference.  
526 *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

527 528 Carl d. Boor. *A Practical Guide to Splines*. Springer Verlag, New York, 1978.

529 530 Maximilian Dax, Stephen R Green, Jonathan Gair, Jakob H Macke, Alessandra Buonanno, and  
531 Bernhard Schölkopf. Real-time gravitational wave science with neural posterior estimation.  
532 *Physical review letters*, 127(24):241103, 2021.

533 534 Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin  
535 loss for deep face recognition. In *Conference on Computer Vision and Pattern Recognition*, 2019.

536 537 Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In  
538 *International Conference on Learning Representations*, 2017.

539 Justin Domke. Provable smoothness guarantees for black-box variational inference. In *International  
Conference on Machine Learning*, 2020.

540 Justin Domke, Robert M. Gower, and Guillaume Garrigos. Provable convergence guarantees for  
541 black-box variational inference. In *Neural Information Processing Systems*, 2023.

540 Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In  
 541 *Advances in Neural Information Processing Systems*, 2019.

542

543 Paul H. C. Eilers and Brian D. Marx. Flexible smoothing with B-splines and penalties. *Statistical*  
 544 *Science*, 11(2):89 – 121, 1996.

545

546 Ankush Ganguly, Sanjana Jain, and Ukrit Watchareeruetai. Amortized variational inference: A  
 547 systematic review. *Journal of Artificial Intelligence Research*, 78:1–49, 1 2024.

548

549 Samuel Gershman, Matt Hoffman, and David Blei. Nonparametric variational inference. In *International*  
 550 *Conference on Machine Learning*, 2012.

551

552 Derek Hansen, Ismael Mendoza, Runjing Liu, Ziteng Pang, Zhe Zhao, Camille Avestruz, and Jeffrey  
 553 Regier. Scalable Bayesian inference for detection and deblending in astronomical images. In *ICML*  
 554 *Workshop on Machine Learning for Astrophysics*, 2022.

555

556 Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *Similarity-Based Pattern*  
 557 *Recognition*, 2015.

558

559 Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and  
 560 generalization in neural networks. In *Neural Information Processing Systems*, 2018.

561

562 Glenn Jocher et al. YOLOv5. <https://github.com/ultralytics/yolov5>, 2020.

563

564 Ilyes Khemakhem, Diederik Kingma, Ricardo Monti, and Aapo Hyvärinen. Variational autoencoders  
 565 and nonlinear ICA: A unifying framework. In *Proceedings of the Twenty Third International*  
 566 *Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research,  
 567 2020.

568

569 Patrick Kidger and Cristian Garcia. Equinox: neural networks in JAX via callable PyTrees and  
 570 filtered transformations. *Differentiable Programming workshop at Neural Information Processing*  
 571 *Systems 2021*, 2021.

572

573 Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Foundations and*  
 574 *Trends in Machine Learning*, 12(4):307–392, 2019.

575

576 J. Lars Kirkby, Álvaro Leitao, and Duy Nguyen. Spline local basis methods for nonparametric density  
 577 estimation. *Statistics Surveys*, 17:75–118, 2023.

578

579 Ja-Yong Koo and Woo-Chul Kim. Wavelet density estimation by approximation of log-densities.  
 580 *Statistics & Probability Letters*, 26(3):271–278, 1996.

581

582 Tuan Anh Le, Adam R. Kosiorek, N. Siddharth, Yee Whye Teh, and Frank Wood. Revisiting  
 583 reweighted wake-sleep for models with stochastic control flow. In *Proceedings of the Thirty-Fifth*  
 584 *Conference on Uncertainty in Artificial Intelligence*, 2019.

585

586 Jongmin Lee, Joo Young Choi, Ernest K Ryu, and Albert No. Neural tangent kernel analysis of deep  
 587 narrow neural networks. In *International Conference on Machine Learning*, 2022.

588

589 Runjing Liu, Jon D. McAuliffe, Jeffrey Regier, and LSST Dark Energy Science Collaboration.  
 590 Variational inference for deblending crowded starfields. *Journal of Machine Learning Research*,  
 591 24(179):1–36, 2023a.

592

593 Runjing Liu, Jon D. McAuliffe, Jeffrey Regier, and LSST Dark Energy Science Collaboration.  
 594 Variational inference for deblending crowded starfields. *Journal of Machine Learning Research*,  
 595 24(179):1–36, 2023b.

596

597 Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. SphereFace: Deep  
 598 hypersphere embedding for face recognition. In *Conference on Computer Vision and Pattern*  
 599 *Recognition*, 2017.

600

601 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Confer-*  
 602 *ence on Learning Representations*, 2019.

594 LSST DESC, Bela Abolfathi, David Alonso, Robert Armstrong, Eric Aubourg, et al. The LSST  
 595 DESC DC2 simulated sky survey. *The Astrophysical Journal Supplement Series*, 253(1):31, 3  
 596 2021.

597 David J. C. MacKay. A practical bayesian framework for backpropagation networks. *Neural  
 598 Computation*, 4(3):448–472, 1992a. doi: 10.1162/neco.1992.4.3.448.

599 David J. C. MacKay. Bayesian interpolation. *Neural Computation*, 4(3):415–447, 05 1992b. ISSN  
 600 0899-7667. doi: 10.1162/neco.1992.4.3.415. URL <https://doi.org/10.1162/neco.1992.4.3.415>.

601 Alex I. Malz and David W. Hogg. How to obtain the redshift distribution from probabilistic redshift  
 602 estimates. *The Astrophysical Journal*, 928(2):127, 2022.

603 Declan McNamara, Jackson Loper, and Jeffrey Regier. Globally convergent variational inference. In  
 604 *Advances in Neural Information Processing Systems*, 2024a.

605 Declan McNamara, Jackson Loper, and Jeffrey Regier. Sequential Monte Carlo for inclusive  
 606 KL minimization in amortized variational inference. In *Proceedings of The 27th International  
 607 Conference on Artificial Intelligence and Statistics*, 2024b.

608 Grant Merz, Xin Liu, Samuel Schmidt, Alex I. Malz, Tianqing Zhang, et al. Deepdisc-photoz: Deep  
 609 learning-based photometric redshift estimation for Rubin LSST. *The Open Journal of Astrophysics*,  
 610 8, 2025.

611 Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg,  
 612 1996. ISBN 0387947248.

613 Art B. Owen. *Monte Carlo Theory, Methods and Examples*, chapter 9, pp. 265–294. Stanford  
 614 University, 2013. Chapter on Importance Sampling. Available at <https://artowen.su.domains/mc/>.

615 Lorenzo Pacchiardi and Ritabrata Dutta. Score matched neural exponential families for likelihood-free  
 616 inference. *Journal of Machine Learning Research*, 23(38):1–71, 2022.

617 George Papamakarios and Iain Murray. Fast  $\epsilon$ -free inference of simulation models with Bayesian  
 618 conditional density estimation. In *Advances in Neural Information Processing Systems*, 2016.

619 George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji  
 620 Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of  
 621 Machine Learning Research*, 22(57):1–64, 2021.

622 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, and  
 623 others Killeen. Pytorch: An imperative style, high-performance deep learning library. In *Advances  
 624 in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.

625 Aakash Patel, Tianqing Zhang, Camille Avestruz, Jeffrey Regier, and The LSST Dark Energy Science  
 626 Collaboration. Neural posterior estimation for cataloging astronomical images with spatially  
 627 varying backgrounds and point spread functions. *The Astronomical Journal*, 170(3):155, 8 2025.

628 Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings  
 629 of the 32nd International Conference on Machine Learning*, 2015.

630 Kihyuk Sohn. Improved deep metric learning with multi-class N-pair loss objective. In *Advances in  
 631 Neural Information Processing Systems*, 2016.

632 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and  
 633 Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th  
 634 International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May  
 635 3-7, 2021*. OpenReview.net, 2021.

636 Bharath Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Aapo Hyvärinen, and Revant Kumar.  
 637 Density estimation in infinite dimensional exponential families. *Journal of Machine Learning  
 638 Research*, 18(57):1–59, 2017.

648 Manoj Kumar Srivastava, Abdul Hamid Khan, and Namita Srivastava. *Statistical Inference: Theory  
649 of Estimation*. PHI Learning, 2014.

650

651 J Michael Steele. *Stochastic calculus and financial applications*. Springer, New York, 2010.

652

653 The JAX Authors. Gpu memory allocation — jax documentation. [https://docs.jax.dev/en/latest/gpu\\_memory\\_allocation.html](https://docs.jax.dev/en/latest/gpu_memory_allocation.html), 2024. Accessed: 2025-11-19.

654

655 Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European  
656 Conference on Computer Vision*, volume 12356 of *Lecture Notes in Computer Science*, pp. 776–  
657 794. Springer, 2020.

658

659 Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational  
660 inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.

661

662 Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei  
663 Liu. Cosface: Large margin cosine loss for deep face recognition. In *2018 IEEE/CVF Conference  
on Computer Vision and Pattern Recognition*, 2018.

664

665 Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R. Scott. Multi-similarity  
666 loss with general pair weighting for deep metric learning. In *Conference on Computer Vision and  
Pattern Recognition*, 2019.

667

668 Daniel Ward, Patrick Cannon, Mark Beaumont, Matteo Fasiolo, and Sebastian Schmon. Robust neural  
669 posterior estimation and statistical model criticism. *Advances in Neural Information Processing  
Systems*, 2022.

670

671 Stephan Wojtowytsch. On the convergence of gradient descent training for two-layer ReLU-networks  
672 in the mean field regime. *arXiv preprint arXiv:2005.13530*, 2020.

673

674 Quan Xiao and Tianyi Chen. Unlocking global optimality in bilevel optimization: A pilot study. In *The Thirteenth International Conference on Learning Representations*, 2025.

675

676 Mang Ye, Xu Zhang, Pong C. Yuen, and Shih-Fu Chang. Unsupervised embedding learning via in-  
677 variant and spreading instance feature. In *Conference on Computer Vision and Pattern Recognition*,  
678 2019.

679

680 Cheng Zhang, Judith Bütepage, Hedvig Kjellström, and Stephan Mandt. Advances in variational  
681 inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):2008–2026,  
682 2019.

683

684 Dingyi Zhang, Yingming Li, and Zhongfei Zhang. Deep metric learning with spherical embedding.  
685 In *Proceedings of the 34th International Conference on Neural Information Processing Systems*,  
2020.

686

687 Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar C Tatikonda, Nicha Dvornek, Xenophon Pa-  
688 pademetris, and James Duncan. Adabelief optimizer: Adapting stepsizes by the belief in observed  
689 gradients. *Advances in Neural Information Processing Systems*, 33, 2020.

690

691

692

693

694

695

696

697

698

699

700

701

## 702 A B-SPLINE & WAVELET BASIS FUNCTIONS 703

704 We give examples of two classes of local basis functions defined on the unit interval  $[0, 1]$ : B-splines  
705 and wavelets. Without loss of generality, these definitions can be extended to construct the family  
706 on any interval  $I = [a, b]$  for  $a, b \in \mathbb{R}$ . These families are potential candidates for a local basis  
707 parameterization of the variational posterior.

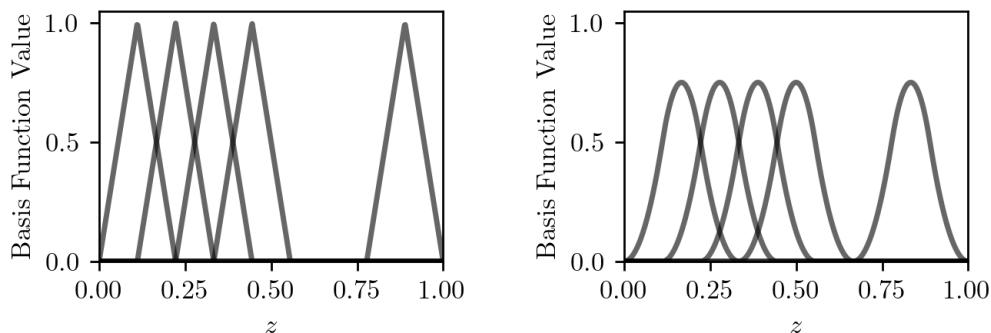
### 709 A.1 B-SPLINES 710

711 For a choice of degree  $d \geq 1$  and a uniformly spaced set of points  $\{t_i\}_{i=1}^K$ , the B-splines are a set of  
712 local basis functions  $\{b_i^{(d)}(z)\}_{i=1}^K$  that are defined recursively as follows (d. Boor, 1978; Eilers &  
713 Marx, 1996):

$$715 \quad b_i^{(d)}(z) = 1 \text{ for all } i, z, \text{ if } d = 0 \quad (13)$$

$$717 \quad \frac{b_i^{(d)}(z)}{(t_{i+d} - t_i)} = \frac{z - t_i}{t_{i+d-1} - t_i} \cdot \frac{b_i^{(d-1)}(z)}{t_{i+d-1} - t_i} + \frac{t_i - z}{t_{i+d} - t_i} \cdot \frac{b_{i+1}^{(d-1)}(z)}{t_{i+d} - t_i} \quad d \geq 1. \quad (14)$$

719 The B-spline basis functions  $b_i^{(d)}(z)$  are thus individually piecewise polynomial splines of degree  $d$ .  
720 While each function  $b_i^{(d)}$  is symmetric about the  $i$ th knot  $t_i$  (or a midpoint of two knots), it is nonzero  
721 for a range of only  $2d$  knots, in accordance with the *locality* aspect described above (see Figure 4).  
722



737 Figure 4: Example visualizations of B-Spline basis functions of degree 1 (left) and 2 (right). For  
738 brevity, we only show a part of the basis functions.

### 740 A.2 WAVELETS 741

742 A collection of wavelet local basis functions on  $[0, 1]$  is defined relative to a “mother wavelet” function  
743 denoted  $H$ . For ease of exposition, we consider the Haar wavelet (Koo & Kim, 1996; Steele, 2010)  
744 given by

$$745 \quad H(z) = \begin{cases} 1, & 0 \leq z < \frac{1}{2} \\ -1, & \frac{1}{2} \leq z \leq 1 \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

748 Thereafter, the set of local basis functions  $b_i$  is defined recursively as follows: writing each  $i$  uniquely  
749 as  $i = 2^j + k$ ,  $j \geq 0$ ,  $0 \leq k < 2^j$ , we have

$$750 \quad b_i(z) = 2^{j/2} H(2^j \cdot z - k). \quad (16)$$

752 The local basis functions are thus defined as shifted and scaled versions of the mother wavelet  $H$ .  
753 One can check that  $b_i(z)$  is nonzero only on the interval  $[k \cdot 2^{-j}, (k + 1) \cdot 2^{-j}]$  for  $i = 2^j + k$ , so  
754 the basis functions become nonzero on increasingly local regions even for moderate values of  $i$  (say,  
755  $i = 200$ ). Wavelets are often described as being local with respect to both space *and* frequency as  $b_i$   
becomes increasingly “spiky” as well due to the coefficient  $2^{j/2}$  (Steele, 2010).

756 **B CONVEXITY**  
 757

758 **B.1 CONVEXITY & CONVERGENCE OF LBF-NPE WITH FIXED BASIS FUNCTIONS**  
 759

760 As referenced in Section 4.2, in the setting where the basis functions  $s_\psi$  are fixed ahead of time, the  
 761 objective function of LBF-NPE becomes a convex functional of the amortization network  $f$ . In this  
 762 case, our setting can be shown to be globally convergent under suitable regularity conditions on the  
 763 network architecture, in the asymptotic limit as the network width tends arbitrarily large. We restate  
 764 this result, proven in McNamara et al. (2024a), below.

765 **Proposition.** *Let  $\mathcal{X} \subseteq \mathbb{R}^d$  and  $\mathcal{Y} \subseteq \mathbb{R}^K$ . Let  $f_\phi = f(\cdot; \phi) : \mathcal{X} \rightarrow \mathbb{R}^K$  be parameterized as a scaled  
 766 two-layer ReLU network of width  $p$ , i.e.  $f_i(x; \phi) = \frac{1}{\sqrt{p}} \sum_{j=1}^p a_{ij} \sigma(x^\top w_j)$  for  $i = 1, \dots, K$ . Define  
 767 the loss functional*

$$768 \quad \ell(x, \eta) = \text{KL}(p(z | x) || q(z; \eta)),$$

769 and allow parameters  $\phi = \{a_{ij}, w_j\}$ ,  $i = 1, \dots, K, j = 1, \dots, p$  to evolve via the gradient flow ODE  
 770  $\dot{\phi}(t) = -\nabla_\phi \mathbb{E}_{p(x)} \ell(x, f(x; \phi(t)))$ . Then we have the following results (under regularity conditions  
 771 (A1)–(A6)):

772 *a)  $L_{\text{LBF-NPE}}(\phi(t))$  is precisely  $\mathbb{E}_{p(x)} \ell(x, f(x; \phi(t)))$ , optimized by the gradient flow above.  
 773 Further, the functional  $M_{\text{LBF-NPE}}(f) : f \mapsto \mathbb{E}_{p(x)} \ell(x, f(x))$  is a convex functional in  $f$ , with  
 774 a global optimum  $f^*$ .*

775 *b) For the parameterization above, with  $\phi$  following the gradient flow ODE, we have that there  
 776 exists  $T > 0$  such that*

$$777 \quad \lim_{p \rightarrow \infty} M_{\text{LBF-NPE}}(f_T) \leq M_{\text{LBF-NPE}}(f^*) + \epsilon,$$

778 where  $f_T = f(\cdot; \phi(T))$ .

779 The proposition above, proven in (McNamara et al., 2024a), states that gradient descent on the convex  
 780 functional  $\mathbb{E}_{p(x)} \ell(x, f(x; \phi))$  converges arbitrarily close to its optimum in the infinite-width limit,  
 781 relying on universality results of shallow networks and NTK theory (Jacot et al., 2018; Lee et al.,  
 782 2022). Our parameterization of the variational distribution  $q$  as an exponential family in LBF-NPE  
 783 falls into this setting when the basis functions are fixed, allowing us to directly apply results from  
 784 (McNamara et al., 2024a). We refer to the proofs therein rather than a restatement here.

785 Regularity conditions sufficient for the above to hold are provided below. They assume a well-behaved  
 786 functional  $M$ , a particular initialization of the width- $p$  network, and uniform boundedness conditions  
 787 on gradients along the optimization trajectory. Although we emphasize that NTK-style results only  
 788 approximately explain the success of our method or other neural posterior estimation methods in  
 789 practice (the infinite-width limit can only be approximated by finite networks, and the continuous  
 790 gradient flow ODE is approximated by stochastic gradient descent), results like these prove that  
 791 LBF-NPE benefits from an advantageous optimization landscape asymptotically.

792 (A1) The data space  $\mathcal{X}$  is compact.

793 (A2) Weights are initialized as  $a_{ij} \stackrel{a.s.}{=} 0$ ,  $w_j \stackrel{iid}{\sim} \mathcal{N}(0, I_d)$ .

794 (A3) The neural tangent kernel at initialization,  $K_\phi(x, x') = Jf(x; \phi) Jf(x'; \phi)^\top$  at  $\phi = \phi(0)$ ,  
 795 is dominated by some integrable random variable  $G$ , uniformly over  $x, x'$ .

796 (A4) The gradient  $\nabla_\eta \ell(x, \eta)$  is uniformly bounded for all  $x, \eta$ .

797 (A5) The limiting NTK  $K_\infty = \lim_{p \rightarrow \infty} K_\phi$  is positive-definite (we note that existence of the  
 798 limit is guaranteed).

799 (A5) The functional  $M_{\text{LBF-NPE}}$  is bounded below; and its minimizer  $f^*$  has finite norm with  
 800 respect to the RKHS norm of the limiting NTK  $K_\infty$ .

801 (A6) The function  $\ell(x, \eta)$  is  $C$ -smooth in  $\eta$  for some  $C < \infty$ .

810 B.2 CONVEXITY OF LBF-NPE  
811

812 In this subsection, we prove the convexity of the joint functional  $L(f, s)$  in Proposition 1. Neural  
813 network theory and NTK-based analysis of this objective in  $s$  are beyond the scope of this work. We  
814 prove marginal convexity in  $s$ , and appeal to previous NTK-based results (such as the above), which  
815 motivate our empirical results: optimization of convex functionals of neural network outputs is ad-  
816 vantageous compared to the optimization of nonconvex functionals due to the preferable optimization  
817 landscape of the former (Bach, 2017; Bengio et al., 2005; Jacot et al., 2018; Wojtowytsch, 2020).

818 We first present Hölder’s inequality, which we’ll use in the proof.  
819

820 **Lemma 1** (Hölder). *If  $S$  is a measurable subset of  $\mathbb{R}^n$  (with respect to Lebesgue measure), and  $f$   
821 and  $g$  are measurable real-valued functions on  $S$ , then Hölder’s inequality is*

$$824 \int_S |f(x)g(x)| dx \leq \left( \int_S |f(x)|^p dx \right)^{\frac{1}{p}} \left( \int_S |g(x)|^q dx \right)^{\frac{1}{q}}.$$

828 for any  $p, q$  satisfying  $\frac{1}{p} + \frac{1}{q} = 1$ .  
829

830 Hölder’s inequality will be used to prove Proposition 1, restated below.  
831

832 **Proposition 1.** *The functional*

$$834 L(f, s) = -\mathbb{E}_{p(z, x)} \left( f(x)^\top s(z) - \log \left( \int \exp(f(x)^\top s(z)) dh(z) \right) \right)$$

837 is marginally convex in the argument  $f$  and  $s$ , respectively.  
838

841 *Proof.* Ignore the outer expectation for now, and consider a fixed  $x, z$  drawn from  $p(z, x)$ . The  
842 inner product  $-f(x)^\top s(z)$  is clearly convex in  $s$ . We now turn to the more complicated expression,  
843  $\log \left( \int \exp(f(x)^\top s(z')) dz' \right)$ . Note that the integral is over  $z'$  this expression and doesn’t depend  
844 on the realization  $z$ . It still depends on  $s$ , however.  
845

846 We prove this function is convex as follows. Let  $\alpha \in (0, 1)$ , and consider functions  $s_1, s_2$ . Then  
847

$$\begin{aligned} 848 & \log \left( \int \exp(f(x)^\top [\alpha s_1(z') + (1 - \alpha)s_2(z')]) dh(z') \right) \\ 849 &= \log \left( \int \exp(\alpha f(x)^\top s_1(z')) \exp((1 - \alpha)f(x)^\top s_2(z')) dh(z') \right) \\ 850 &= \log \left( \int [\exp(f(x)^\top s_1(z'))]^\alpha [\exp(f(x)^\top s_2(z'))]^{1-\alpha} dh(z') \right) \\ 851 &= \log \left( \int [u(z')]^\alpha [v(z')]^{1-\alpha} dh(z') \right) \end{aligned}$$

858  
859  
860 where we have defined  $u(z') = \exp(f(x)^\top s_1(z'))$  and  $v(z') = \exp(f(x)^\top s_2(z'))$ . Consider the  
861 integral above, and apply Hölder’s inequality with  $1/p = \alpha$  and  $1/q = 1 - \alpha$ . These sum to one as  
862 required. We take  $f$  to be  $[u(z')]^\alpha$  and  $g$  to be  $[v(z')]^{1-\alpha}$ .  
863

Continuing, we have

$$\begin{aligned}
&= \log \left( \int [u(z')]^\alpha \cdot [v(z')]^{(1-\alpha)} dh(z') \right) \quad [\text{this line repeated for clarity}] \\
&\leq \log \left( \left[ \int ([u(z')]^\alpha)^{1/\alpha} dh(z') \right]^\alpha \cdot \left[ \int ([v(z')]^{(1-\alpha)})^{1/(1-\alpha)} dh(z') \right]^{1-\alpha} \right) \quad [\text{Hölder}] \\
&= \alpha \log \left( \int u(z') dh(z') \right) + (1-\alpha) \log \left( \int v(z') dh(z') \right) \\
&= \alpha \log \left( \int \exp(f(x)^\top s_1(z')) dh(z') \right) + (1-\alpha) \left( \int \exp(f(x)^\top s_2(z')) dh(z') \right).
\end{aligned}$$

This was all that is required to show that the mapping  $s \mapsto \log \left( \int \exp(f(x)^\top s(z')) dz' \right)$  is convex. As the sum of two convex functions is convex, we've shown convexity of the integrand for any fixed draw  $z, x \sim p(z, x)$ . To conclude, we observe that by linearity of integration, this holds for the integral as well.

□

## C SAMPLING

Similar to EigenVI (Cai et al., 2024), LBF-NPE does not easily admit sampling from the fitted variational density. This is one limitation of the nonparametric nature of the density model in both of these approaches to variational inference.

In the low-dimensional case, sampling can be performed by inverse transform sampling. In this approach, one uses the cumulative distribution function  $Q$  of  $q$ , defined as

$$Q(z^*) = P_q(z < z^*) = \int_{-\infty}^{z^*} q(z) dz$$

where  $q(z)$  is the fitted variational density (say, conditional on some datum  $x$  of interest). The function  $Q$  is invertible. Sampling can be performed by drawing  $U \sim \text{Unif}[0, 1]$ , and thereafter computing

$$Z = Q^{-1}(U).$$

The result of this procedure is a draw  $Z \sim q(z)$ . Computing and inverting cumulative distribution functions in low dimensions is fairly straightforward. As outlined in the main text, this low-dimensional setting is one we commonly find to be of use to practitioners, especially for the types of problems we consider in our experiments in this work.

To sample from high-dimensional posteriors, several different approaches are available. One approach, as outlined in (Cai et al., 2024), is sequential sampling, whereby one samples

$$q(z_1), q(z_2 | z_1), q(z_3 | z_1, z_2), \dots, q(z_d | z_1, \dots, z_{d-1})$$

in order. Each individual density above can be sampled using the inverse transform sampling approach outlined above; conditioning within the exponential family parameterization is accomplished easily by freezing already sampled indices and changing the variables of integration in the log integral. More generally, one could also utilize rejection sampling or other Monte Carlo sampling algorithms. As the unnormalized variational density has an extremely simple form, i.e.  $q(z) \propto \exp(\eta^\top b(z))$ , Markov chain Monte Carlo algorithms could also be an efficient way to sample from the distribution defined by the fitted density. We present some selected results of Langevin sampling and inverse transform sampling in Appendix C.1 to illustrate the utility of these approaches.

### C.1 RESULTS OF LANGEVIN DYNAMICS AND INVERSE TRANSFORM SAMPLING

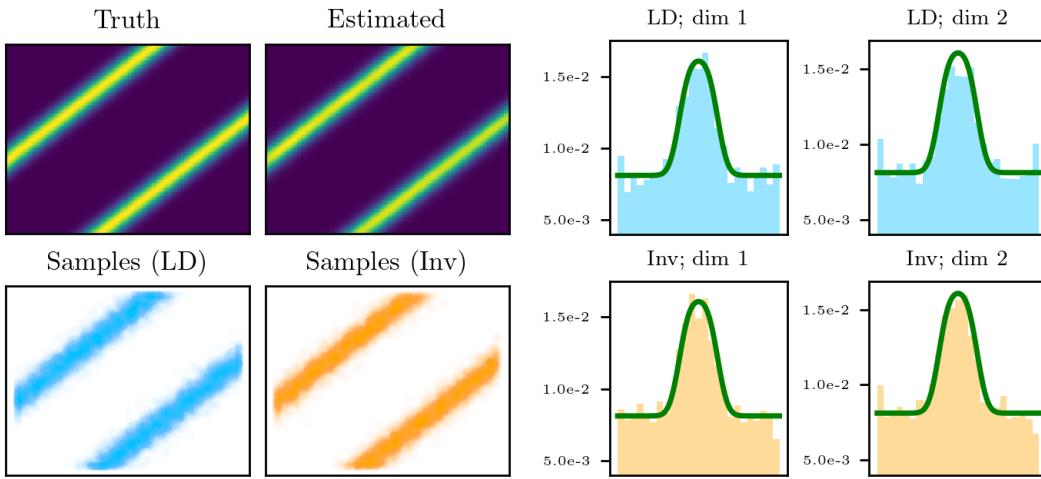
We present the sampling results obtained via Langevin dynamics (Song et al., 2021) and inverse transform sampling for the three 2D case studies: `bands`, `ring`, and `spiral`. In addition to the inverse transform sampling described in the preceding section, we explore the use of Langevin

918 dynamics, a widely adopted method for generating samples in score-based generative models. This  
 919 approach iteratively updates a set of particles according to:  
 920

$$dz_t = \epsilon \nabla_z \log p(z_t | x) dt + \sqrt{2\epsilon} dW_t,$$

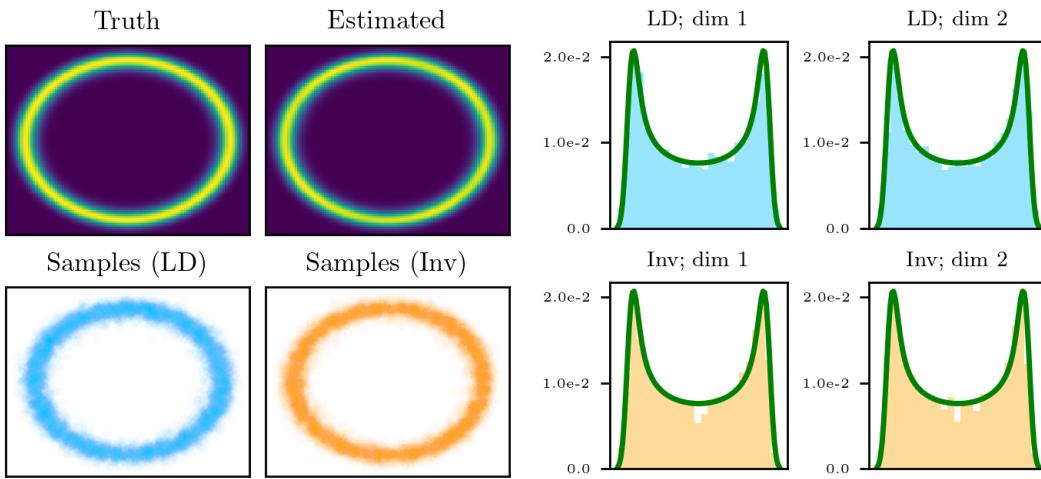
924 where  $\epsilon = 0.001$ ,  $dW_t \sim \mathcal{N}(0, 1)$ , and  $t \in \{1, 2, \dots, 1000\}$ . Since our model provides a differentiable  
 925 approximation of  $\log p(z_t | x)$  through  $f_\phi(x)^\top s_\psi(z)$ , the gradient  $\nabla_z \log p(z_t | x)$  can be  
 926 directly estimated. This enables us to apply Langevin dynamics for posterior sampling.

927 For each of the bands, ring, and spiral case studies, we generate 10,000 samples and visualize  
 928 both the samples and their marginal densities in Figures 5 to 10. As illustrated in these figures, both  
 929 Langevin dynamics and inverse transform sampling yield samples that closely match the estimated  
 930 posterior distributions.



948 Figure 5: Sampling results for bands. “LD” and  
 949 “Inv” denote Langevin dynamics and inverse trans-  
 950 form sampling, respectively.

Figure 6: Marginal density of samples for bands.  
 The green line indicates the estimated posterior  
 density.



969 Figure 7: Sampling results for ring. “LD” and  
 970 “Inv” denote Langevin dynamics and inverse trans-  
 971 form sampling, respectively.

Figure 8: Marginal density of samples for ring.  
 The green line indicates the estimated posterior  
 density.

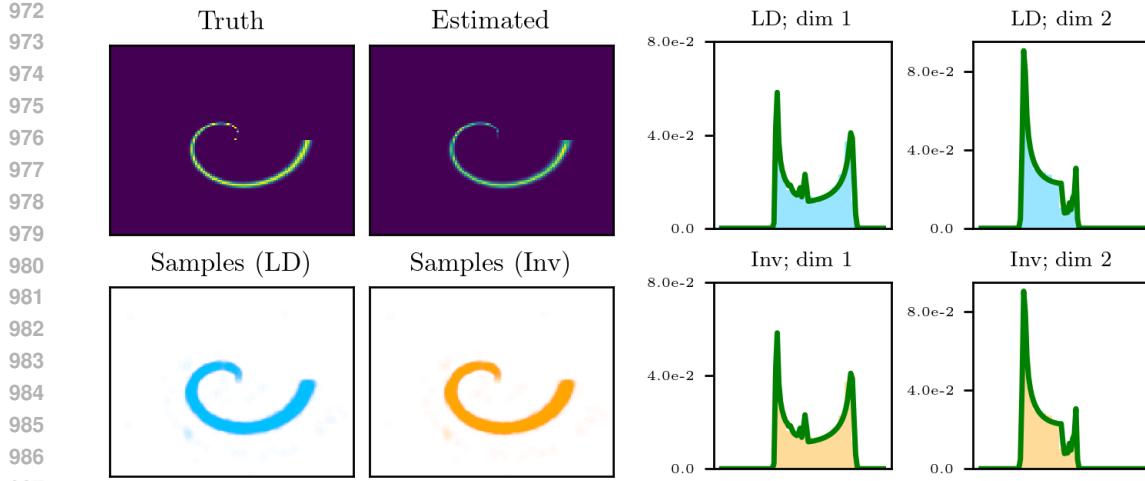


Figure 9: Sampling results for spiral. “LD” and “Inv” denote Langevin dynamics and inverse transform sampling, respectively.

Figure 10: Marginal density of samples for spiral. The green line indicates the estimated posterior density.

## D EXPERIMENTAL DETAILS

Code to reproduce results is provided at [https://anonymous.4open.science/r/lbf\\_npe-725A](https://anonymous.4open.science/r/lbf_npe-725A) (to be de-anonymized). We use PyTorch (Paszke et al., 2019; Ansel et al., 2024) and JAX (Bradbury et al., 2018). We also use the `equinox` library for deep learning in JAX (Kidger & Garcia, 2021). For Section 6.4, the DC2 Simulated Sky Survey data is publicly available from the LSST Dark Energy Science Collaboration (LSST DESC) (LSST DESC et al., 2021; Collaboration et al., 2022). All experiments are conducted on an Ubuntu 22.04 server equipped with an NVIDIA RTX 2080 Ti GPU.

### D.1 TOY EXAMPLE: SINUSOIDAL LIKELIHOOD

In Section 6.1, we design a simple Bayesian model to evaluate the convergence behavior of LBF-NPE and MDN. This statistical model is

$$\begin{aligned} z &\sim \text{Unif}[0, 2\pi], \\ x \mid z &\sim \mathcal{N}(\sin(2z), \sigma^2), \end{aligned}$$

where  $\sigma^2 = 1$ . This model induces a multimodal posterior:

$$P(z \mid x) \propto \exp\left(-(\sin(2z) - x)^2 / (2\sigma^2)\right) \times \mathbb{I}(0 \leq z \leq 2\pi),$$

which exhibits two modes when  $x \geq 1$  or  $x \leq -1$ , and four modes otherwise.

For LBF-NPE, we construct a multilayer perceptron (MLP) to predict the coefficient vector  $\eta = f_\phi(x)$ , and the sufficient statistics  $b(z) \in \mathbb{R}^K$  are computed using B-spline basis functions with  $K = 14$ . The MLP architecture consists of an input layer, four hidden layers, and an output layer, mapping  $x \in \mathbb{R}$  to  $\eta \in \mathbb{R}^K$ . Each hidden layer includes a full connection layer of 128 units, layer normalization, and a ReLU activation. The output layer is linear. The B-spline basis comprises 14 degree-2 basis functions, with knots at  $[0, 0, \text{linspace}(0, 2\pi, \text{num} = K), 2\pi, 2\pi]$ . Although the B-spline basis  $b(z)$  can be evaluated recursively as described in Section A.1, we precompute it on a grid to avoid redundant computation during training. Specifically, we pick 1000 uniformly spaced points in the interval  $[0, 2\pi]$  and approximate the integral  $\int \exp(\eta_i^\top b(z)) dz$  using the trapezoidal sum. For the term  $f_\phi(x)^\top b(z)$ , we use the basis vector corresponding to the grid point closest to the true latent variable  $z_{\text{true}}$ .

For MDN, we use the same MLP architecture, with output adapted to represent the parameters of a mixture of Gaussian distributions with  $K = 14$  components. The output vector has 10 parameters for means and variances, and 4 additional parameters for the mixture weights. The MDN objective is

$$L_{\text{MDN}}(\gamma) = -\mathbb{E}_{p(x,z)} \log q(z; t_\gamma(x)), \quad (17)$$

1026 where  $\gamma$  are the neural network parameters,  $t_\gamma(x)$  denotes the predicted distribution parameters, and  
 1027  $q(z; \cdot)$  is the corresponding density.  
 1028

1029 The training procedures for LBF-NPE and MDN are identical apart from the loss function. At  
 1030 each step, we sample 1024 latent–observed pairs  $(z, x)$  from the generative model and update model  
 1031 parameters using the AdaBelief optimizer (Zhuang et al., 2020) with a learning rate of 0.001. Training  
 1032 proceeds for 50,000 steps and completes within one hour for both methods. Peak GPU memory usage  
 1033 is approximately 8300MB. We hold out 1000  $(z, x)$  pairs and track their negative log-likelihood  
 1034 (NLL) over training, as shown in Figure 1. We apply Gaussian smoothing to the NLL curves with  
 1035 standard deviation  $\sigma = 20$ . This results in a smoothing kernel of size  $161 = 4 \times 20 \times 2 + 1$ , with  
 1036 weights given by  $G_i = \exp(-i^2/(2\sigma^2))$  for  $i \in \{-80, \dots, 80\}$ . With the normalization constants  
 1037 for  $G_i$  omitted, the smoothed NLL at step  $j$  is computed as  
 1038

$$1039 \text{NLL}_{\text{smooth},j} = \sum_{i=-80}^{80} \text{NLL}_{j+i} \cdot G_i.$$

1040  
 1041 **D.2 2D CASE STUDIES: BANDS**  
 1042

1043 The statistical model for the bands test case, as introduced in Section 6.2, is  
 1044

$$1045 z_1, z_2 \sim \text{Unif}[-1, 1], \\ 1046 z = (z_1, z_2), \\ 1047 x \mid z \sim \mathcal{N}(|z_1 - z_2|, \sigma^2),$$

1048 where  $\sigma^2 = 10^{-2}$ . The resulting posterior forms two elongated bands in the 2D latent space  
 1049  $P(z \mid x) \propto \exp(-(|z_1 - z_2| - x)^2/(2\sigma^2)) \cdot \mathbb{I}(-1 \leq z_1, z_2 \leq 1)$ , with its maxima occurring along  
 1050 the lines where  $|z_1 - z_2| = x$ .  
 1051

1052 As the latent variable  $z$  is now two-dimensional, LBF-NPE encounters increased complexity due  
 1053 to the larger number of basis functions required. In our LBF-NPE framework, both the coefficient  
 1054 network  $f_\phi$  and the sufficient statistic network  $s_\psi$  are implemented as multilayer perceptrons (MLPs)  
 1055 with four hidden layers, each containing 128 units. All layers use layer normalization to stabilize  
 1056 optimization and ReLU activations. The network  $f_\phi$  maps input  $x \in \mathbb{R}$  to a coefficient vector in  
 1057  $\mathbb{R}^K$ , while  $s_\psi$  maps  $z \in \mathbb{R}^2$  to sufficient statistics in  $\mathbb{R}^K$ . We set  $K = 20$  for consistency with  
 1058 other 2D case studies, though even  $K = 2$  suffices to capture the posterior structure in this example  
 1059 (see Appendix E.3). The loss function for LBF-NPE follows Algorithm 1, where the integral term  
 1060  $\int \exp(\eta_i^\top s_\psi(z)) dz$  is approximated using a trapezoidal sum over a  $100 \times 100$  uniform grid spanning  
 1061  $[-1, 1]^2$ . During training, we alternate between updating  $f_\phi$  and  $s_\psi$ : we train  $f_\phi$  for 1000 steps  
 1062 while holding  $s_\psi$  fixed, then train  $s_\psi$  for 1000 steps with  $f_\phi$  fixed, and repeat this process until the  
 1063 total training budget is exhausted. The choice of 1000 steps per phase is empirical; we observe  
 1064 diminishing returns in the loss reduction beyond 1000 steps, indicating that each sub-network has  
 1065 reached a near-optimal solution given the other is fixed. In addition, we use stereographic projection  
 1066 to reparameterize the output of  $f_\phi$  and  $s_\psi$ .

1067 For the MDN baseline, we use an MLP with the same architecture as  $f_\phi$ , except that it outputs a  
 1068 50-dimensional vector representing the parameters of a mixture of 10 Gaussian components. Each  
 1069 component is parameterized by five values: two for the mean, two for the (diagonal) variance  
 1070 (assuming zero covariance), and one for the mixture weight. The loss function is identical to that  
 1071 described in Appendix D.1.

1072 For the normalizing flow baseline, we adapt the classic coupling flow from (Dinh et al., 2017) to  
 1073 model the conditional posterior  $p(z \mid x)$ . Each coupling layer includes translation and scaling  
 1074 sub-networks that are conditioned on  $x \in \mathbb{R}$ . These sub-networks are implemented as MLPs, each  
 1075 taking as input the concatenation of the masked latent variable  $z$  and the conditioning variable  $x$ .  
 1076 Each MLP consists of a single hidden layer with 128 units. We use 10 coupling layers to ensure  
 1077 sufficient expressiveness. The resulting conditional density is given by:  
 1078

$$q(z \mid x) = q_{\mathcal{N}}(h_\nu(z; x)) \cdot |\det J|,$$

1079 where  $q_{\mathcal{N}}(\cdot)$  denotes the standard Gaussian density,  $h_\nu(z; x)$  is the transformed latent variable via  
 1080 the flow, and  $\det J$  is the product of the Jacobian determinants from each flow transformation.

We train LBF-NPE, MDN, and the normalizing flow using the AdamW optimizer (Loshchilov & Hutter, 2019) with a learning rate of  $10^{-5}$  for 50,000 steps. The batch size is set to 1024, and training completes in approximately 2 hours. Maximum GPU memory usage is around 8400MB. For evaluation, we use a held-out set of 1000  $(z, x)$  pairs to compute the average forward and reverse KL divergences. For each test observation  $x$ , LBF-NPE, MDN, and the normalizing flow estimate the density  $q(z | x)$  over a  $100 \times 100$  uniform grid on  $[-1, 1]^2$ . These estimated posteriors are normalized such that their integral over the grid equals 1. The true posterior  $p(z | x)$  is computed analytically over the same grid, enabling pointwise comparison. We then calculate the forward and reverse KL divergences between the estimated and true posteriors and report the average over all 1000 test cases in Table 1. For the illustrative posterior plots shown in Figure 2, we fix  $x = 0.7$  and visualize the estimated density  $q(z | x)$  from each method over the same  $100 \times 100$  grid.

### D.3 2D CASE STUDIES: RING

The statistical model for the `ring` case study in Section 6.2 is defined as:

$$\begin{aligned} z_1, z_2 &\sim \text{Unif}[-1, 1], \\ z &= (z_1, z_2), \\ x | z &\sim \mathcal{N}(\|z\|^2, \sigma^2), \end{aligned}$$

where  $\sigma^2 = 10^{-2}$ . The resulting posterior,  $P(z | x) \propto \exp\left(-(\|z\|^2 - x)^2/(2\sigma^2)\right) \cdot \mathbb{I}(-1 \leq z_1, z_2 \leq 1)$ , forms a ring-shaped distribution in the latent space, with radius approximately  $\sqrt{x}$ .

The network architectures and training configurations used in this case are identical to those described in Appendix D.2. An example posterior  $q(z | x = 0.7)$  is visualized in Figure 2.

### D.4 2D CASE STUDIES: SPIRAL

The `spiral` model is defined as follows:

$$\begin{aligned} b &\sim \text{Unif}[0.1, 0.5] \\ d &\sim \text{Unif}[0.0, s_b(2\pi)] \\ \theta &= s_b^{-1}(d) \\ r &= b\theta \\ z_1 &= r \cos(\theta) \\ z_2 &= r \sin(\theta) \\ z &= (z_1, z_2) \\ x | z &\sim \mathcal{N}(b, \sigma^2) \end{aligned}$$

where  $\sigma^2 = 10^{-4}$ , and  $s_b(\theta) = \frac{b}{2}(\theta\sqrt{1+\theta^2} + \sinh^{-1}(\theta))$ . The posterior is  $P(z | x) \propto \exp\left(-(\frac{r}{\theta} - x)^2/(2\sigma^2)\right) \cdot \mathbb{I}(0 \leq \theta \leq 2\pi, 0.1\theta \leq r \leq 0.5\theta)$ .

Most training settings follow those in Appendix D.2, except that we increase the number of coupling layers in normalizing flow to 16. We observe minimal performance gain beyond this depth. For visualization in Figure 2, we display the estimated posterior  $q(z | x = 0.35)$  over the  $100 \times 100$  grid.

### D.5 OBJECT DETECTION

We define the image generative model as follows:

$$\begin{aligned} l_1, l_2 &\sim \text{Unif}([0, 16] \times [0, 16]), \\ f_1, f_2 &\sim \mathcal{N}(\mu, \sigma^2), \\ I &= \text{Image}(\{l_1, l_2\}, \{f_1, f_2\}, \text{PSF}), \\ x_{j,k} &\sim \text{Poisson}(I_{j,k}), \end{aligned}$$

1134 where  $\mu = 2000$ ,  $\sigma^2 = 400^2$ , and  $\text{Image}(\cdot)$  and  $\text{PSF}(\cdot)$  are defined below. Note that in our  
 1135 implementation, flux values are constrained to be positive.  
 1136

---

1137 **Algorithm 2:** `Image`

1139 **Inputs:** list of source locations  $\{l_1, l_2\}$ ; list  
 1140 of source fluxes  $\{f_1, f_2\}$ ;  
 1141 point-spread function  $\text{PSF}$ .  
 1142 Initialize pixel location matrix  $pl$   
 1143 **for**  $l_i, f_i$  in  $\text{zip}(\{l_1, l_2\}, \{f_1, f_2\})$  **do**  
 1144   | Compute relative location  $rl_i = pl - l_i$   
 1145   | Compute PSF density  $d_i = \text{PSF}(rl_i)$   
 1146   | Compute  $I_i = f_i \times d_i$   
 1147 **end**  
 1148 Compute  $I = I_1 + I_2$   
 1149 Return  $I$ .

---

1150  
 1151 The pixel location matrix  $pl$  is a mesh grid of shape  $(H, W, 2)$  defined over  $[0.5, 1.5, \dots, H - 0.5] \times$   
 1152  $[0.5, 1.5, \dots, W - 0.5]$ , where  $H$  and  $W$  are the height and width of the image, respectively. Each  
 1153 source location  $l_i$  is a 2D vector, and the term  $rl_i[\dots, 0]^2 + rl_i[\dots, 1]^2$  is a matrix of shape  $(H, W)$ .  
 1154 We use  $\sigma_{\text{PSF}}^2 = 1$ . Each source flux  $f_i$  is a scalar. Before passing the image to the network, we  
 1155 normalize it using min-max scaling:  $x = (x - \min(x)) / (\max(x) - \min(x))$ .

1156 Since the input is a  $16 \times 16$  image, we employ a convolutional layer in our network to reduce  
 1157 computational cost. The first layer of the model  $f_\phi$  is a 2D convolution layer with kernel size 4,  
 1158 increasing the channel dimension from 1 to 3. This is followed by a 2D max pooling layer and a  
 1159 ReLU activation. The output is then flattened and passed through four MLP layers, each with 128  
 1160 hidden units, layer normalization, and ReLU activations. Another model,  $s_\psi$ , is an MLP with four  
 1161 hidden layers, each also consisting of 128 units, layer normalization, and ReLU. The outputs of  $f_\phi$   
 1162 and  $s_\psi$  are reparameterized via stereographic projection.

1163 As each image contains two astronomical sources, we compute the loss separately for each source.  
 1164 For the source located at  $l_i$ , the first term in the loss is  $-w\hat{f}_\phi(x)^\top \hat{s}_\psi(l_i)$ . Only  $\hat{s}_\psi(l_i)$  needs to be  
 1165 evaluated per source; shared terms such as  $\hat{f}_\phi(x)$  and the integral term can be reused across both. We  
 1166 approximate the integral using Monte Carlo integration with 22,500 random samples. The final loss is  
 1167 the sum of the losses for both sources. For alternating optimization, we train one of  $f_\phi$  or  $s_\psi$  for 300  
 1168 steps at a time (shorter than the 1000-step updates used in the 2D case studies; see Appendix D.2)  
 1169 since convergence is typically achieved more rapidly in this setting. Optimization is performed using  
 1170 the AdamW optimizer (Loshchilov & Hutter, 2019) with a learning rate of 0.001. The total number  
 1171 of training steps is 45,000, with overall training time under two hours.

1172 For posterior visualization, we adopt the same procedure used in previous 2D case studies (see  
 1173 Appendix D.2), but evaluate the posterior over a  $200 \times 200$  grid on the domain  $[0, 16]^2$ . The estimated  
 1174 posterior for a certain image is shown in Figure 3. To generate this posterior, we leverage a model  
 1175 trained with  $K = 64$ . For results with other values of  $K$  (e.g., 9, 20, 36), we provide further  
 1176 discussion in Appendix E.3.

1177  
 1178 **D.6 REDSHIFT ESTIMATION**

1180 Our redshift experiment extends the methodology of the Bayesian Light Source Separator (BLISS)  
 1181 (Liu et al., 2023a; Hansen et al., 2022; Patel et al., 2025). For a given generative model of astronomical  
 1182 images and latent quantities (locations; fluxes; type of object; redshift; etc.), BLISS utilizes neural  
 1183 posterior estimation (Papamakarios & Murray, 2016) to perform amortized variational inference. The  
 1184 network architecture for BLISS operates on *tiles* of images, returning distributional parameters for  
 1185 each object detected per tile. The architecture is thus convolutional in nature with several additional  
 1186 image normalizations and other design choices suitable for astronomical image processing.

1187 For samples of the generative model, we use images from two tracts of the LSST DESC DC2  
 1188 Simulated Sky Survey (LSST DESC et al., 2021; Collaboration et al., 2022), numbers 3828 and 3829.

1188 LBF-NPE does not sample the generative model on-the-fly in this setting, but only have access to a  
 1189 finite number of draws from the training sets.  
 1190

1191 We utilize the BLISS preprocessing routines to produce training, validation, and test sets of images  
 1192 along with ground-truth catalogs. Images, each of size  $80 \times 80$ , are processed in batches of 64 by the  
 1193 BLISS inference network, which further splits these into  $4 \times 4$ -pixel tiles. The network is fit to the  
 1194 training set to minimize the forward KL divergence using a learning rate of 0.001. All nuisance latent  
 1195 variables are marginalized over, and we only score redshift variational posteriors, although BLISS  
 1196 allows for easy addition of posteriors on other latent quantities in the computed NLL loss as well,  
 1197 should the user desire to perform inference on these.  
 1198

1199 We adapt the neural network architecture from BLISS for redshift estimation. The complete architec-  
 1200 tural details and parameter configurations are provided in Table 3. As shown in the table, the input  
 1201 and output shapes of each layer are expressed as tuples, e.g.,  $(\text{bands}, h, w)$  or  $(64, h, w)$ ,  
 1202 where bands denotes the number of bands in the input astronomical images. In the DC2 dataset,  
 1203 there are six bands:  $u, g, r, i, z, y$ . The variables  $h$  and  $w$  represent the height and width of the image,  
 1204 which are both set to 80 in our experiments. The model is composed primarily of three types of  
 1205 layers: Conv2DBlock, C3Block, and Upsample. A Conv2DBlock is a composite module  
 1206 consisting of a 2D convolution, group normalization, and a SiLU activation function. The C3Block  
 1207 is adapted from the YOLOv5 architecture (Jocher et al., 2020). It comprises three convolutional  
 1208 layers with kernel size 1 and includes skip connections implemented via multiple bottleneck blocks  
 1209 (parameterized by  $n$ ). The Upsample layer performs spatial upscaling of the input tensor by a  
 1210 specified factor. The architecture follows a U-shaped design with four downsampling steps followed  
 1211 by two upsampling steps. To denote skip connections and input dependencies between layers, we  
 1212 use the “Input From” column. For example, the entry “[ $-1, 9$ ]” indicates that the current layer takes  
 1213 as input the concatenation of the output from the previous layer and that from layer 9. The final  
 1214 layer is a convolutional module with kernel size 1, producing an output of shape  $(n_{\text{coeff}}, h/4, w/4)$ ,  
 1215 where  $n_{\text{coeff}}$  is the number of predicted coefficients per tile.  
 1216

1217 The forward KL divergence framework prescribes that predictions are only scored for true objects.  
 1218 Accordingly, for each ground-truth redshift in the training set, we score the predicted NLL computed  
 1219 from the variational distribution for the  $4 \times 4$  pixel tile containing that object. BLISS makes this  
 1220 transdimensional inference problem (a result of the number of objects per-tile being unknown *a*  
 1221 *priori*) tractable by sharing parameters among objects within the same  $4 \times 4$ -pixel tile, at the cost of  
 1222 the bias resulting from this approximation. For both the MDN and B-spline parameterization, we fit  
 1223 to the training data for 30 epochs, and use the model weights that had the lowest held-out NLL on the  
 1224 validation set to compute the test-set NLL. Training the inference network  $f_{\phi}$  takes approximately  
 1225 12 hours on a single NVIDIA GeForce RTX 2080 Ti GPU. We note that due to the approximations  
 1226 involved with using a finite training set rather than true “simulated” draws, we can easily overfit to  
 1227 the training and validation set. The procedure outlined above aims to mitigate these issues to the  
 1228 extent possible.  
 1229

1230  
 1231  
 1232  
 1233  
 1234  
 1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241

1242					
1243					
1244					
1245					
1246					
1247					
1248					
1249					
1250					
1251					
1252					
Layer #	Input From	Input Shape	Layer Type	Config	Output Shape
1253					
1254	1	-1	(bands, h, w)	Conv2DBlock	in_ch=bands; out_ch=64; kernel_size=5
1255					(64, h, w)
1256	2	-1	(64, h, w)	Conv2DBlock	in_ch=64; out_ch=64; kernel_size=5
1257					(64, h, w)
1258	3	-1	(64, h, w)	Sequence of Conv2DBlock	in_ch=64; out_ch=64; kernel_size=5; sequence_len=3
1259					(64, h, w)
1260	4	-1	(64, h, w)	Conv2DBlock	in_ch=64; out_ch=64; kernel_size=3; stride=2
1261					(64, h/2, w/2)
1262	5	-1	(64, h/2, w/2)	C3Block	in_ch=64; out_ch=64; n=3
1263					(64, h/2, w/2)
1264	6	-1	(64, h/2, w/2)	Conv2DBlock	in_ch=64; out_ch=128; kernel_size=3; stride=2
1265					(128, h/4, w/4)
1266	7	-1	(128, h/4, w/4)	C3Block	in_ch=128; out_ch=128; n=3
1267					(128, h/4, w/4)
1268	8	-1	(128, h/4, h/4)	Conv2DBlock	in_ch=128; out_ch=256; kernel_size=3; stride=2
1269					(256, h/8, w/8)
1270	9	-1	(256, h/8, w/8)	C3Block	in_ch=256; out_ch=256; n=3
1271					(256, h/8, w/8)
1272	10	-1	(256, h/8, w/8)	Conv2DBlock	in_ch=256; out_ch=512; kernel_size=3; stride=2
1273					(512, h/16, w/16)
1274	11	-1	(512, h/16, w/16)	C3Block	in_ch=512; out_ch=256; n=3
1275					(256, h/16, w/16)
1276	12	-1	(256, h/16, w/16)	Upsample	scale=2; mode="nearest"
1277					(256, h/8, w/8)
1278	13	[-1, 9]	(512, h/8, w/8)	C3Block	in_ch=512; out_ch=256; n=3
1279					(256, h/8, w/8)
1280	14	-1	(256, h/8, w/8)	Upsample	scale=2; mode="nearest"
1281					(256, h/4, w/4)
1282	15	[-1, 6]	(384, h/4, w/4)	C3Block	in_ch=384; out_ch=256; n=3
1283					(256, h/4, w/4)
1284	16	-1	(256, h/4, w/4)	Conv2D	in_ch=256; out_ch=n_coeff; kernel_size=1
1285					(n_coeff, h/4, w/4)

Table 3: Neural network architecture for redshift estimation.

1286  
1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295

1296 D.7 ANGULAR DISTANCE OPTIMIZATION  
1297

1298 As discussed in Section 4.4, our method can be interpreted as performing angular distance optimization,  
1299 but with loss and gradient derived from a probabilistic space. This interpretation becomes evident  
1300 if we decouple the magnitude and directional components of the output tensors  $f_\phi(x), s_\psi(z) \in \mathbb{R}^K$   
1301 through normalization techniques such as L2 normalization or stereographic projection reparameteri-  
1302 zation. Angular distance optimization is a common objective in modern machine learning pipelines,  
1303 contributing to both improved performance and consistent alignment between training and testing  
1304 metrics. Several widely-used loss functions, including the triplet loss (Hoffer & Ailon, 2015), N-pair  
1305 loss (Sohn, 2016), and multi-similarity loss (Wang et al., 2019), incorporate angular distance in their  
1306 formulation. Cosine-based softmax loss functions are employed extensively for face recognition (Liu  
1307 et al., 2017; Wang et al., 2018; Deng et al., 2019), and many contrastive learning algorithms (Chen  
1308 et al., 2020; Tian et al., 2020; Ye et al., 2019) utilize angular objectives to maximize the cosine  
1309 similarity between embeddings from positive pairs.

1310 Our variational objective in Equation (6) is related to cosine-based softmax loss, whose general form  
1311 is

$$1312 L = -wS_{i,y_i} + \log \left( \exp(wS_{i,y_i}) + \sum_{j \neq y_i} \exp(wS_{i,j}) \right), \quad (18)$$

1314 as described in Section 4.4, and suggesting a more general form for our variational objective, namely  
1315

$$1317 \hat{L}_{\text{LBF-NPE}}(\phi, \psi) = \mathbb{E}_{p(z,x)} \left[ -w\hat{f}_\phi(x)^\top \hat{s}_\psi(z) + \log \left( \int \exp \left( w\hat{f}_\phi(x)^\top \hat{s}_\psi(z') \right) dz' \right) \right], \quad (19)$$

1320 where again  $\hat{f}_\phi(\cdot)$  and  $\hat{s}_\psi(\cdot)$  are normalized outputs of neural networks (i.e., with unit norm) and  $w$   
1321 is a scaling factor.

1322 The key differences from the cosine-based softmax formulation are: (1) the summation is replaced by  
1323 an integral over the continuous latent space, and (2) the angular distance is computed between coeffi-  
1324 cient vectors and basis functions, rather than between learned embeddings. This connection offers  
1325 two main advantages. First, our theoretical guarantees on convexity and convergence may extend to  
1326 angular distance optimization problems, suggesting broader applicability. Second, our method can  
1327 leverage off-the-shelf improvements developed for angular optimization, such as SEC (Zhang et al.,  
1328 2020), which regularizes gradient updates to stabilize and accelerate training. Given that even simple  
1329 stereographic normalization already yields strong empirical results, we leave the integration of these  
1330 enhancements to future work.

1331 In our experiments, we utilize stereographic projection reparameterization to normalize the output  
1332 tensor onto the unit hypersphere. Precisely,  $u \in \mathbb{R}^{K-1}$  is transformed to  $\mathbb{R}^K$  via

$$1334 y = \left( \frac{2u}{1 + \|u\|^2}, \frac{1 - \|u\|^2}{1 + \|u\|^2} \right), \quad (20)$$

1336 ensuring that  $\|y\| = 1$ . This projection serves as a smooth and bijective transformation from Euclidean  
1337 space  $\mathbb{R}^{K-1}$  onto the  $K$ -sphere  $S^K = \{x \in \mathbb{R}^K : \|x\| = 1\}$ . Although this transformation changes  
1338 the form of the variational objective in the neural network outputs, and thus violates some assumptions  
1339 of our NTK framing from the perspective of convexity, strong empirical results suggest the benefits  
1340 of reparameterization, and also the importance of future work in understanding the success of neural  
1341 posterior estimation (NPE) techniques under arbitrary parameterizations. We hypothesize that the  
1342 advantageous properties of this normalization stem from the smooth gradient trajectories and mapping  
1343 to a compact space, discussed in more detail below.

1344 We illustrate the stereographic normalization process in a 2D case, as shown in Figure 11. In this  
1345 setting, a scalar input  $u \in \mathbb{R}^1$  is projected onto a vector  $y \in \mathbb{R}^2$  lying on the 1-sphere (i.e., the unit  
1346 circle). For any given  $u$ , there exists a unique line connecting the point  $(u, 0)$  and the north pole  
1347  $N = (0, 1)$ . This line intersects the 1-sphere at a single point, which serves as the projection of  $u$ . By  
1348 drawing a vector from the origin to this intersection point, we obtain a unit vector  $y$  on the 1-sphere.  
1349 Notably, the location of the intersection reflects the magnitude of  $u$ : if  $\|u\| > 1$ , the intersection lies  
on the upper half of the circle; if  $\|u\| < 1$ , it falls on the lower half.

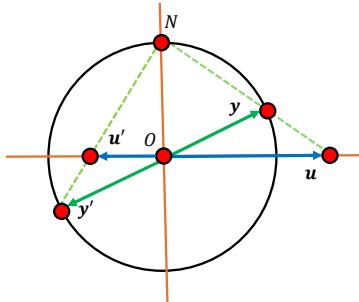


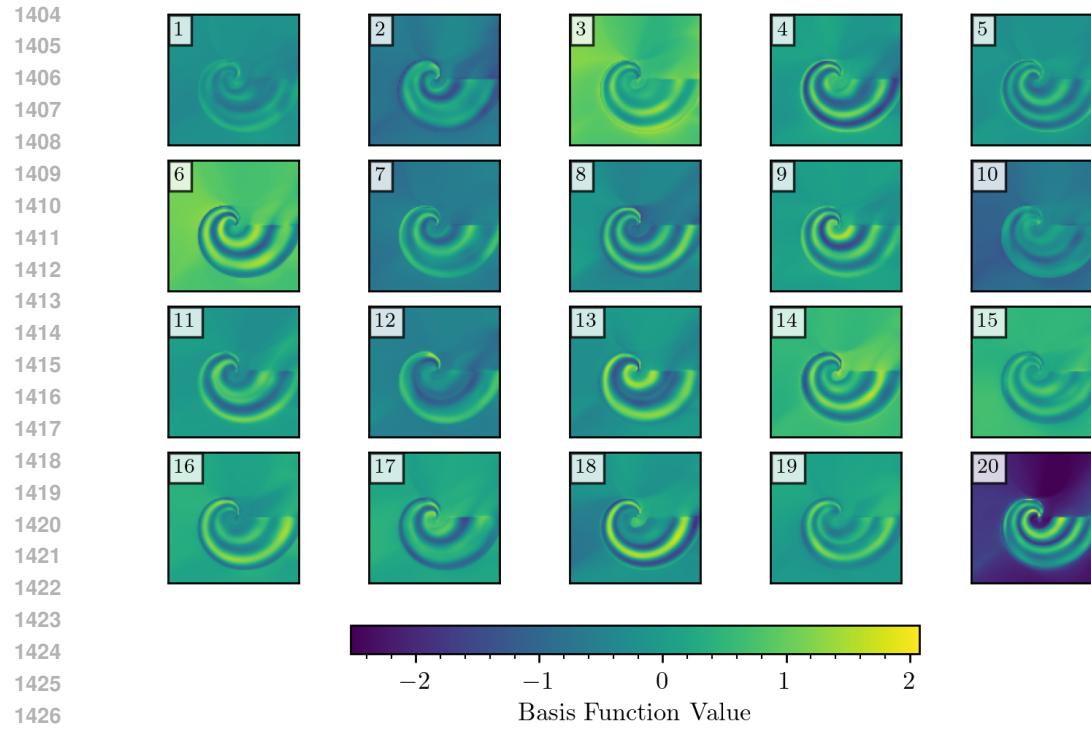
Figure 11: Visualization of stereographic projection in 2D. A scalar  $u$  is mapped to a point on the unit circle via intersection of the line connecting  $(u, 0)$  and the north pole  $N = (0, 1)$ .

This reparameterization offers several advantages. First, the stereographic projection is differentiable everywhere and provides well-behaved gradients throughout the domain. Second, the projection naturally enforces unit-norm constraints without requiring additional normalization layers or manual clipping, thus making training more stable and efficient.

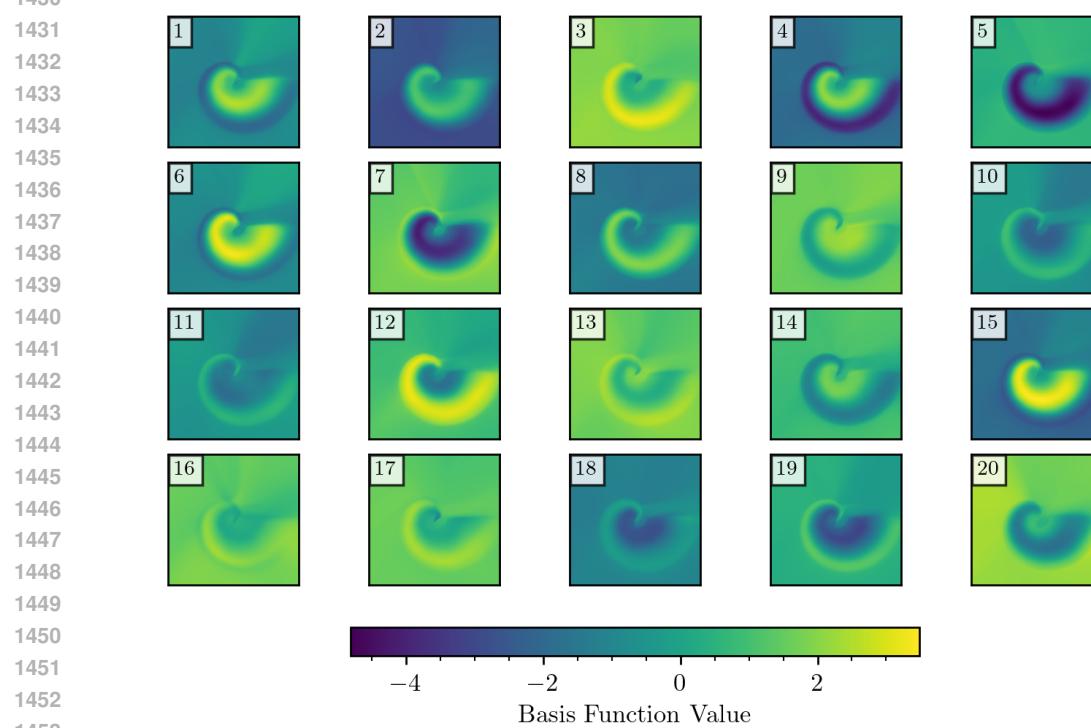
## E ADDITIONAL EXPERIMENTAL RESULTS

### E.1 EFFECT OF NORMALIZATION

We compare the output basis functions of the neural network with and without stereographic projection normalization, demonstrating that normalization helps the network learn clearer boundaries between different regions of the parameter space. Figure 12 and Figure 13 show the values of the 20-dimensional basis functions (i.e.,  $[s_\psi(z)]_i$ , where  $i \in \{1, 2, \dots, 20\}$ ) evaluated over the plane  $z \in [-3, 3] \times [-3, 3]$  for the spiral case study. It is evident that the model with normalization exhibits more distinguishable and structured partition boundaries in  $z$ -space, while the model without normalization suffers from blurry transitions and over-exposure artifacts, as seen in Figure 13. This highlights a key drawback of the non-normalized approach that it struggles to disentangle the parameter space effectively. Normalization also enhances interpretability by promoting better separation among basis functions. The estimated posterior density is expressed as a weighted linear combination of these basis function densities. For a given target spiral, the neural network increases the weights (i.e.,  $[f_\phi(x)]_i$ ) for dimensions whose corresponding basis functions have high overlap with the target density, and decreases weights for dimensions with low overlap.



1428 Figure 12: Density plot of 20-dim basis function (w/ stereographic projection normalization) over  
1429 plane  $z$ . Each subplot represents the density plot of a certain dimension.



1455 Figure 13: Density plot of 20-dim basis function (w/o stereographic projection normalization) over  
1456 plane  $z$ . Each subplot represents the density plot of a certain dimension.

1457

1458  
1459

## E.2 POSTERIOR COMPARISON

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

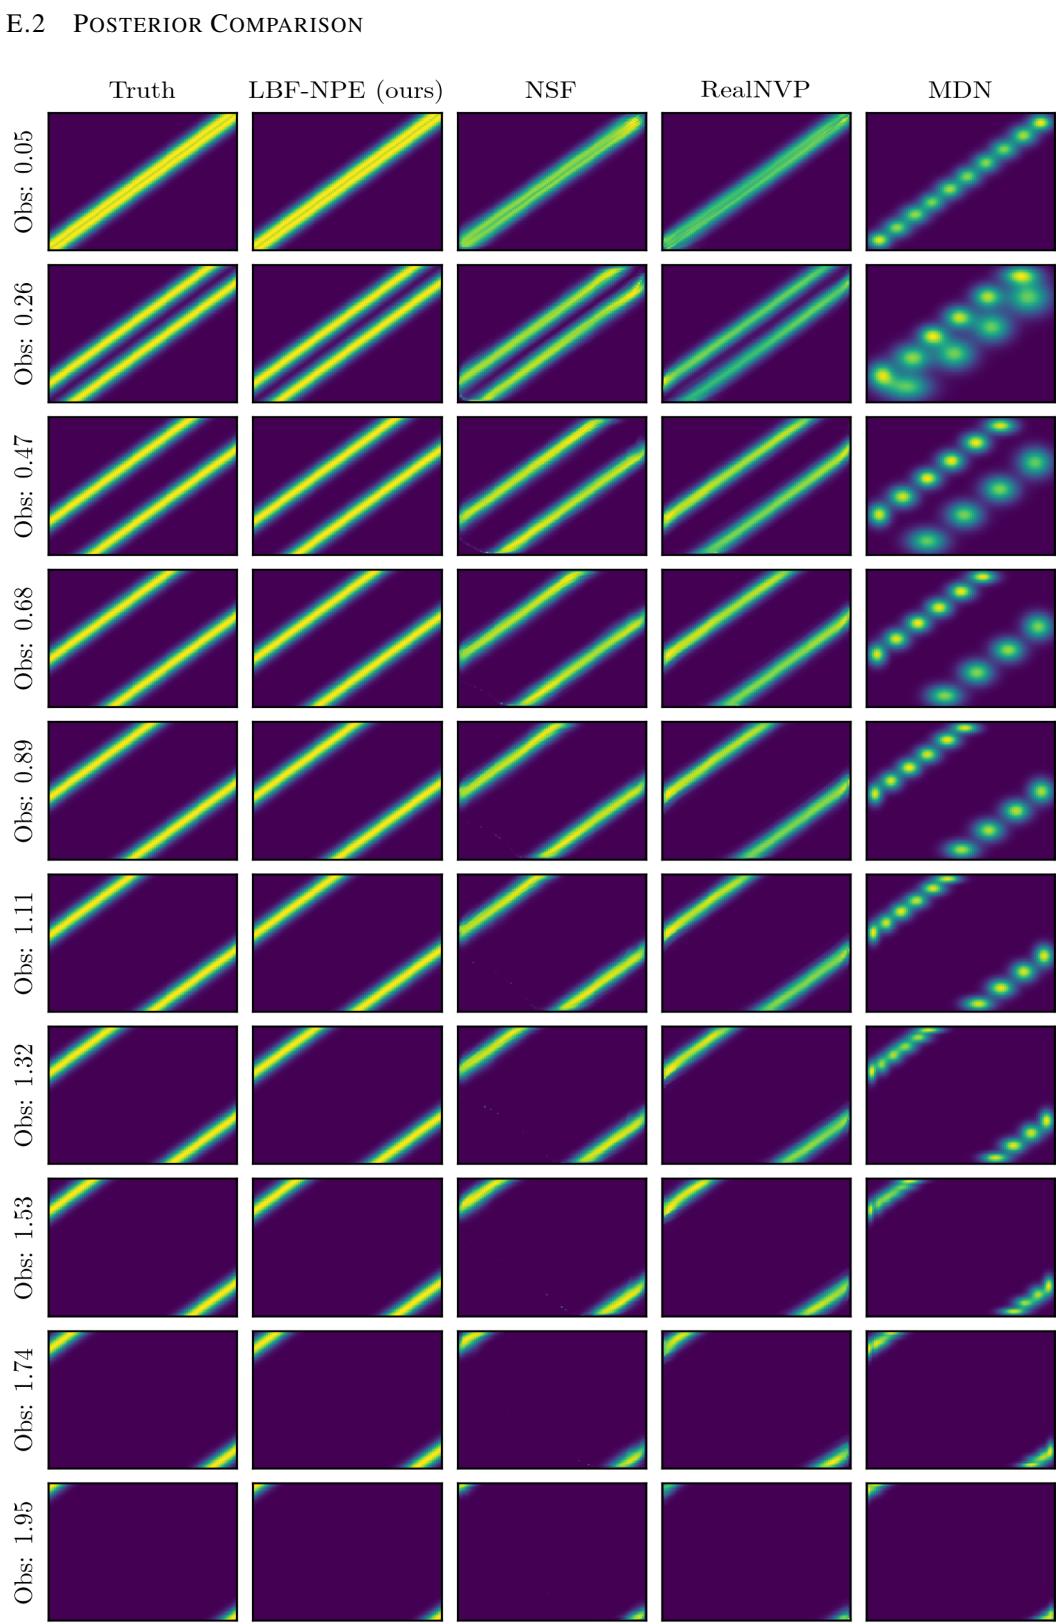


Figure 14: Visual comparison of posterior densities estimated by different methods (LBF-NPE, Normalizing Flow, MDN) against ground truth for ten representative observations.

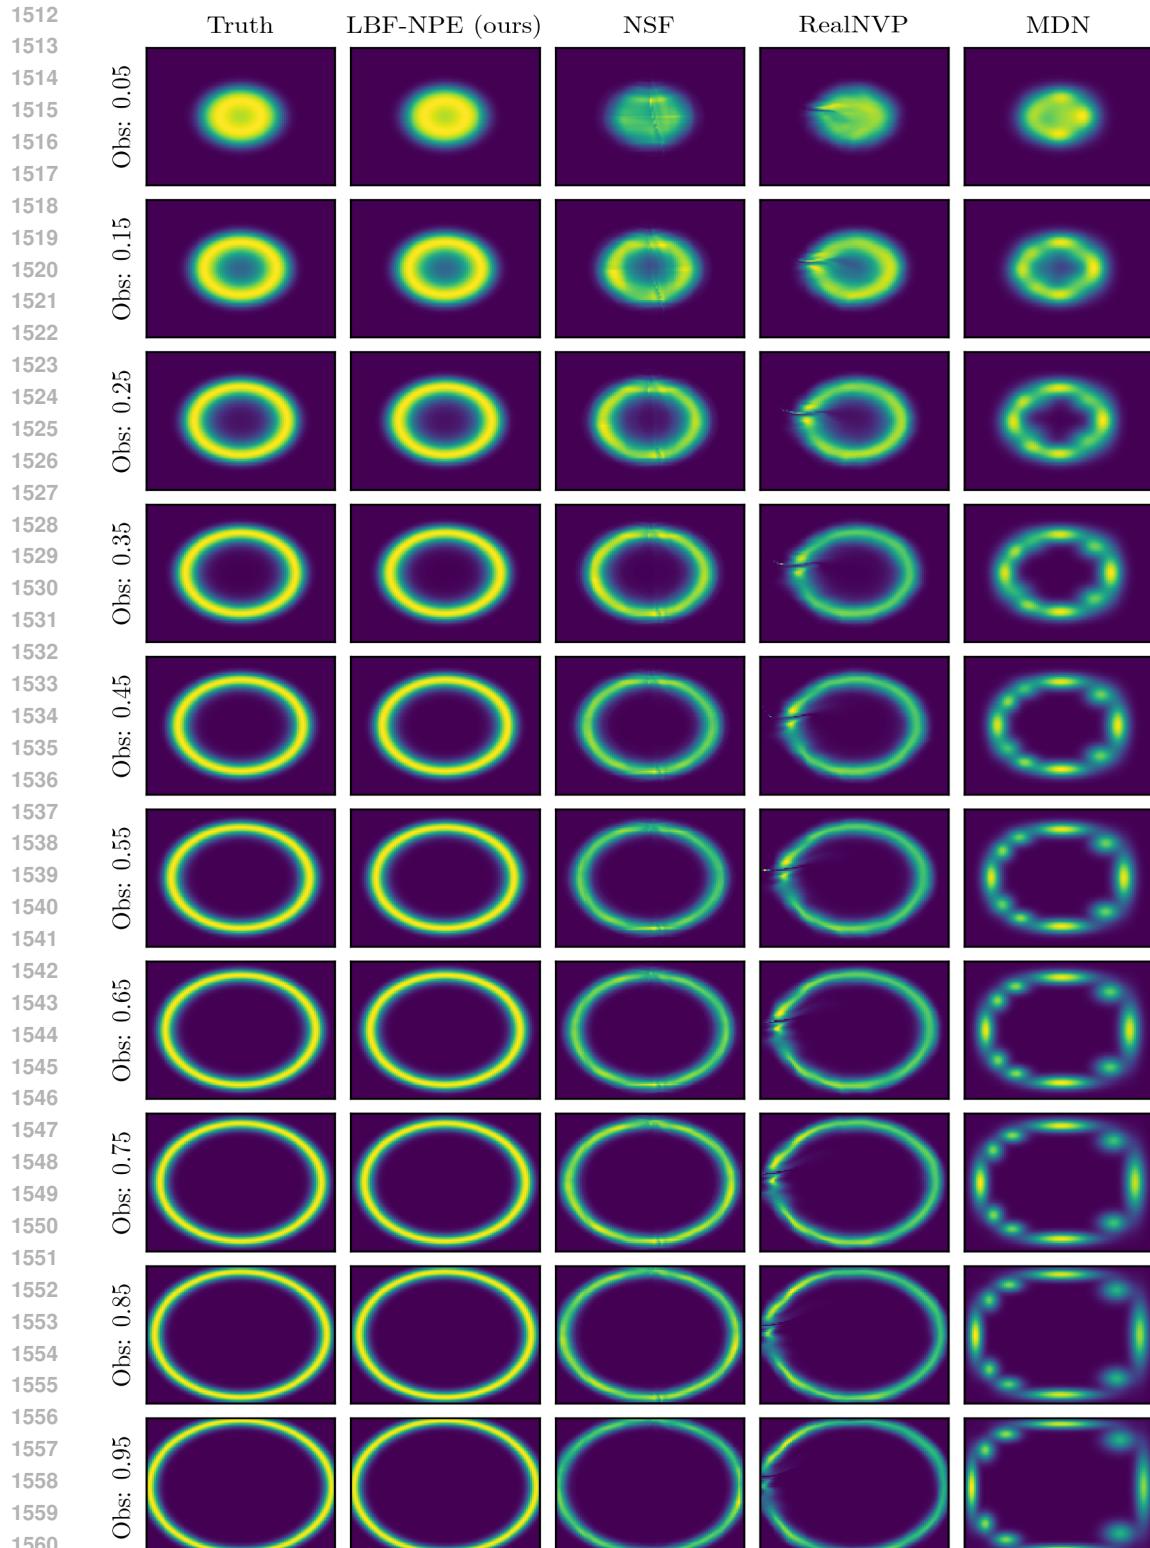


Figure 15: Visual comparison of posterior densities estimated by different methods (LBF-NPE, Normalizing Flow, MDN) against ground truth for ten representative observations.

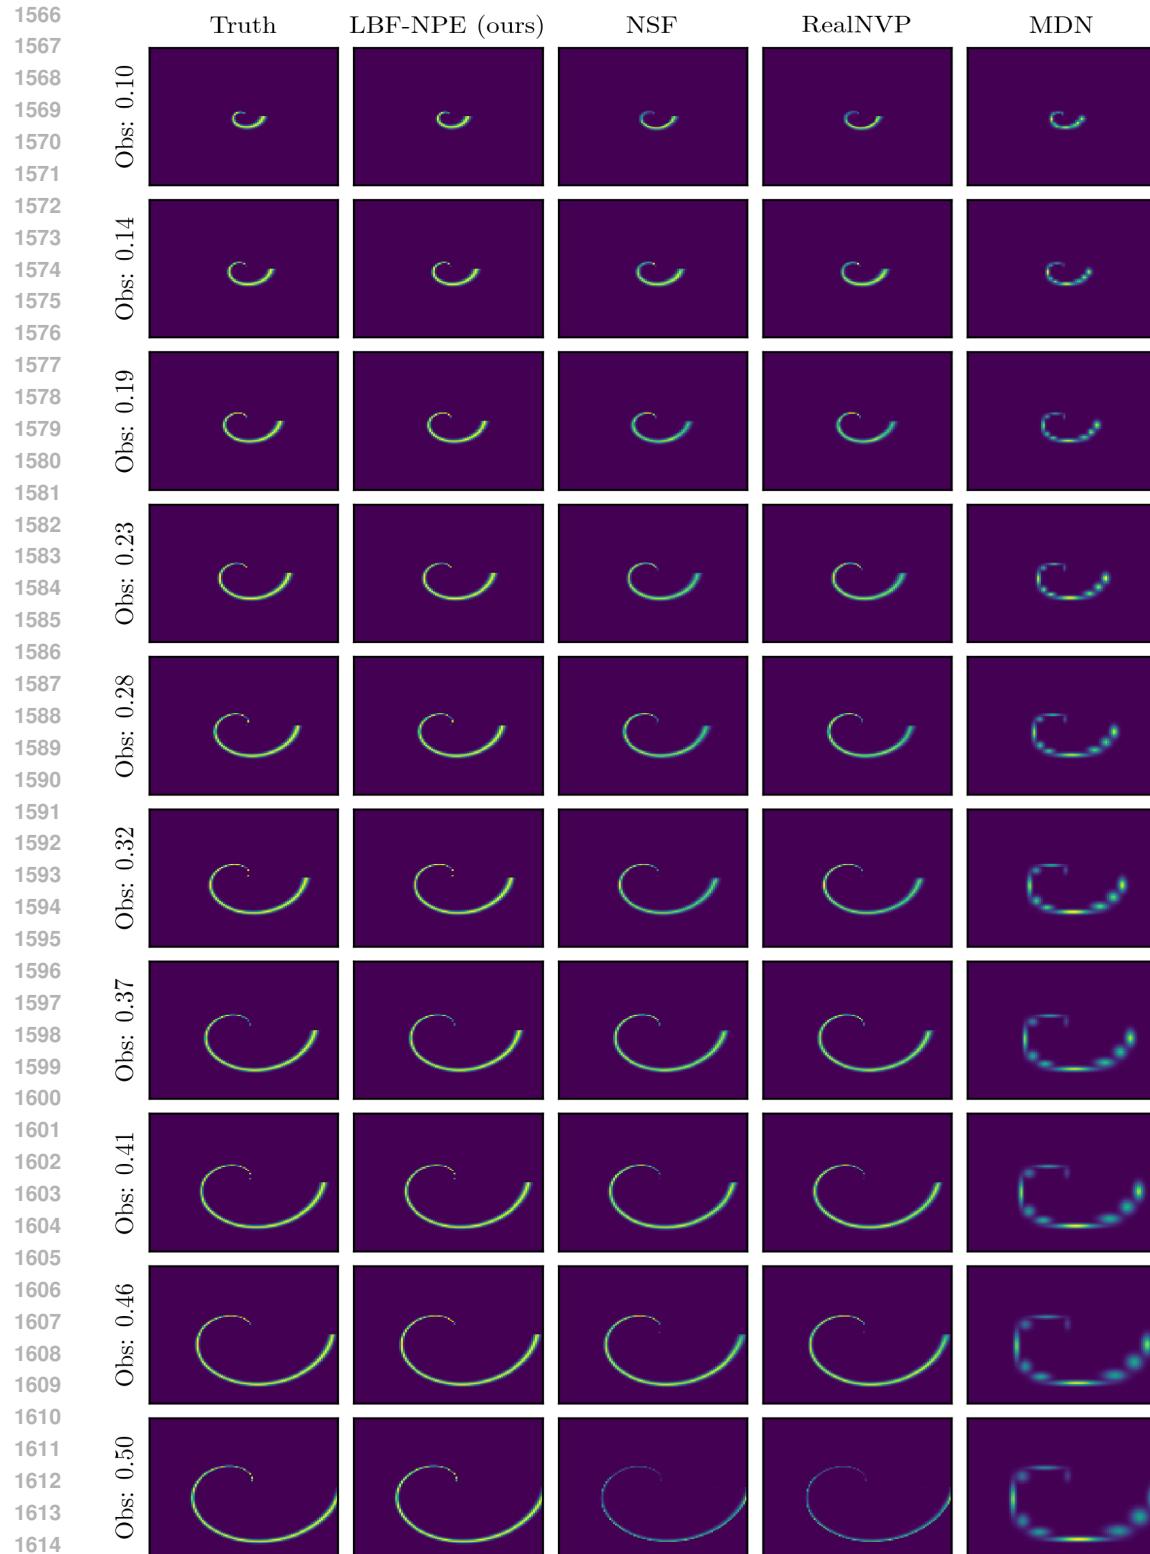


Figure 16: Visual comparison of posterior densities estimated by different methods (LBF-NPE, Normalizing Flow, MDN) against ground truth for ten representative observations.

1620  
1621

## E.3 DIMENSIONS OF BASIS FUNCTIONS &amp; FLEXIBILITY

1622  
1623  
1624  
1625  
1626

The flexibility of our method is positively correlated with the dimensionality of the basis functions. We demonstrate this by analyzing both the forward and reverse KL divergences, as well as the basis function density plots, for models using basis functions of varying dimensions in the object detection case study. To quantify this, we compute the forward and reverse KL divergence between the estimated posterior distribution and a target mixture of Gaussians:

1627  
1628  
1629

$$0.5\mathcal{N}(\text{true locs of source 1}, 0.1^2) + 0.5\mathcal{N}(\text{true locs of source 2}, 0.1^2). \quad (21)$$

1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637

The results, shown in Table 4, indicate that both forward and reverse KL divergences decrease as the dimensionality of the basis functions increases. For example, the 64-dimensional basis functions achieve the lowest forward KL divergence (1.524), while the 9-dimensional basis functions result in the highest (3.397). A similar trend holds for reverse KL divergence. However, the marginal gain from increasing dimensionality diminishes as the number of basis functions grows. Increasing from 9 to 20 dimensions yields a significant improvement in forward/reverse KL divergence (1.187/0.511), but the improvement from 36 to 64 dimensions is relatively minor (0.246/0.287). This suggests that, for a complex task like object detection, a basis function dimensionality under 100 is sufficient to achieve near-optimal performance.

1638  
1639  
1640  
1641  
1642  
1643  
1644

	9-dim	20-dim	36-dim	64-dim
<b>Forward KL Divergence</b>	3.397	2.210	1.770	<b>1.524</b>
$\Delta$ <b>Forward KL Divergence</b>	-	<b>-1.187</b>	-0.440	-0.246
<b>Reverse KL Divergence</b>	2.380	1.869	1.360	<b>1.073</b>
$\Delta$ <b>Reverse KL Divergence</b>	-	<b>-0.511</b>	-0.509	-0.287

1645  
1646  
1647

Table 4: Object detection: forward/reverse KL divergence for models of different basis function dimensions.

1648  
1649  
1650  
1651  
1652

The basis function density plots provide further intuition for this trend. As shown in Figures 17 to 20, the 64- and 36-dimensional basis functions can partition the  $z$ -space into fine-grained regions, capturing detailed structure. In contrast, 20- and 9-dimensional basis functions fail to do so, resulting in coarser approximations and reduced representational capacity.

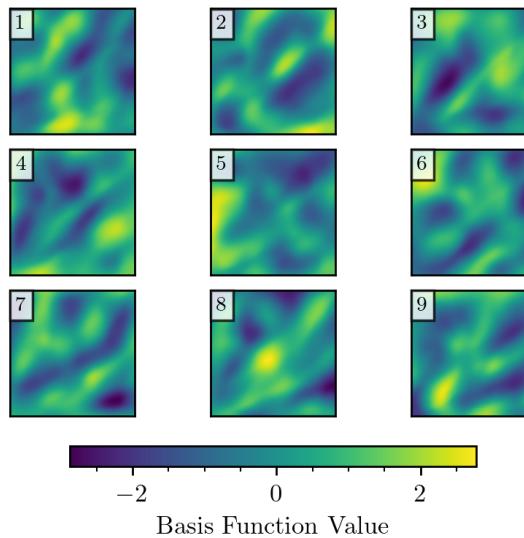
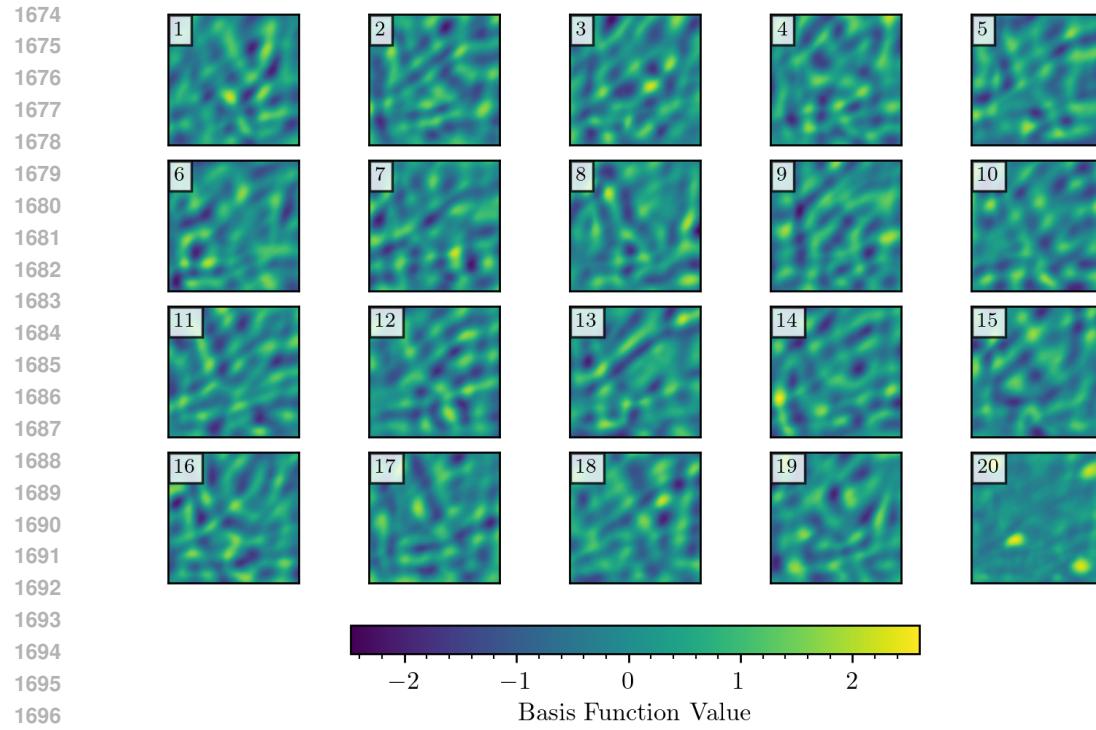
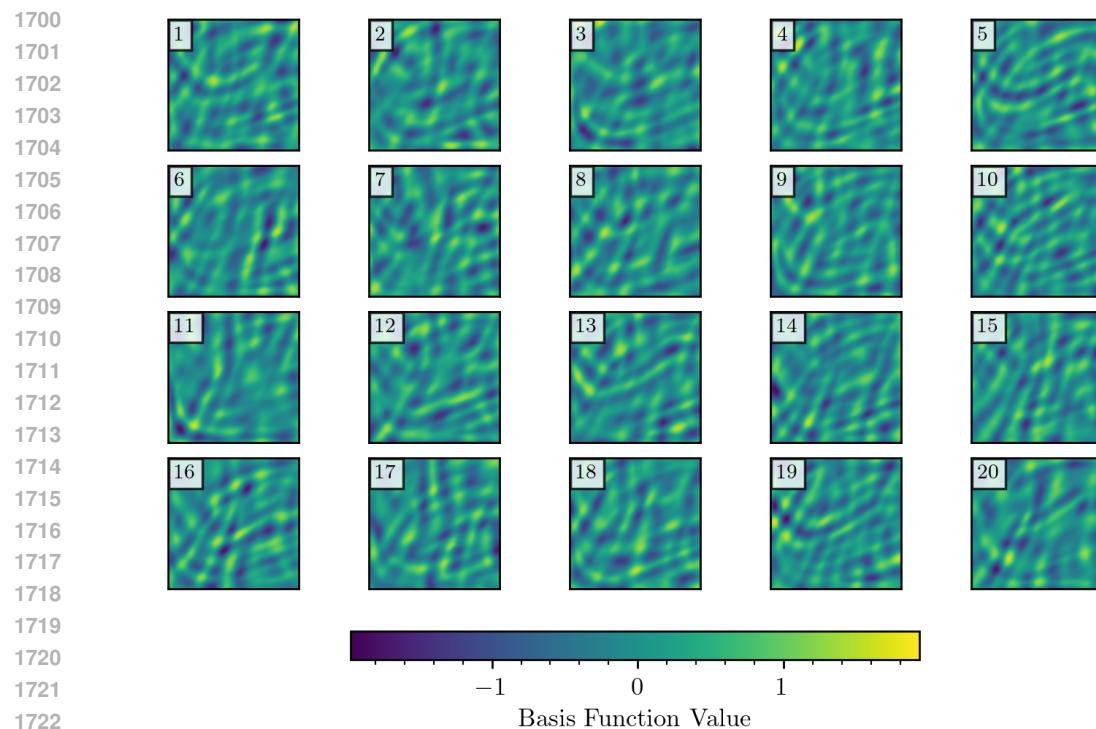
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
16701671  
1672  
1673

Figure 17: Object detection: density plot of 9-dim basis function over plane  $z$ . Each subplot represents the density plot of a certain dimension.

Figure 18: Object detection: density plot of 20-dim basis function over plane  $z$ .Figure 19: Object detection: density plot of 36-dim basis function over plane  $z$ . For brevity, we only show the first 20 dimensions

1724  
1725  
1726  
1727

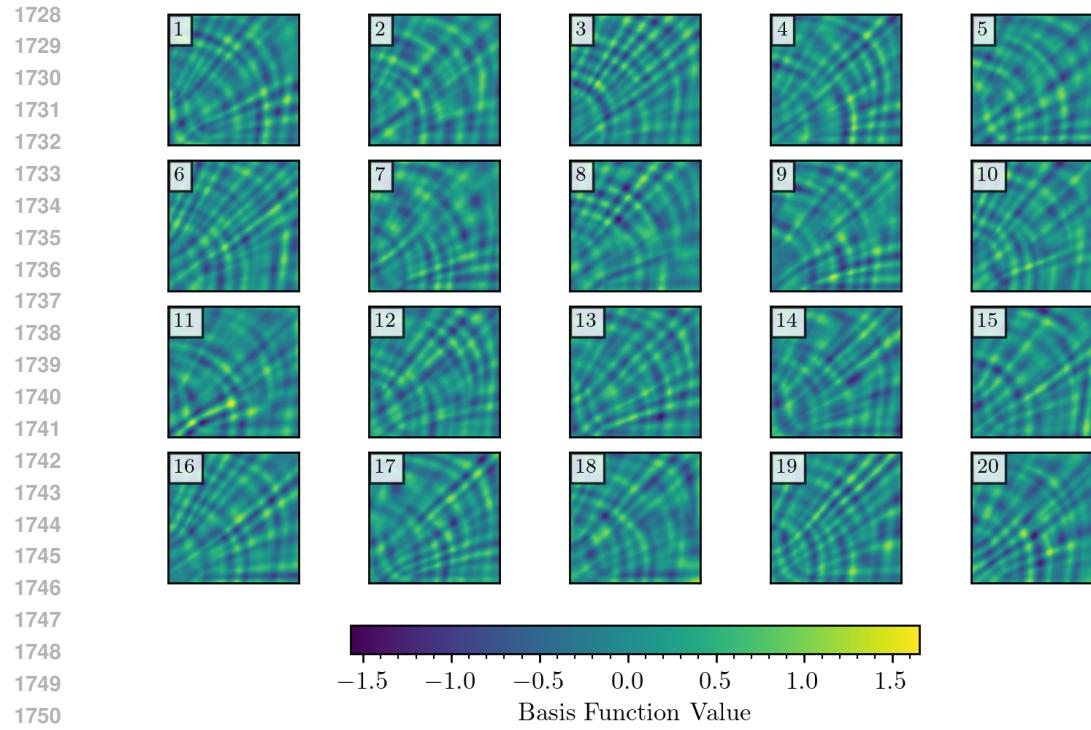


Figure 20: Object detection: density plot of 64-dim basis function over plane  $z$ . For brevity, we only show the first 20 dimensions

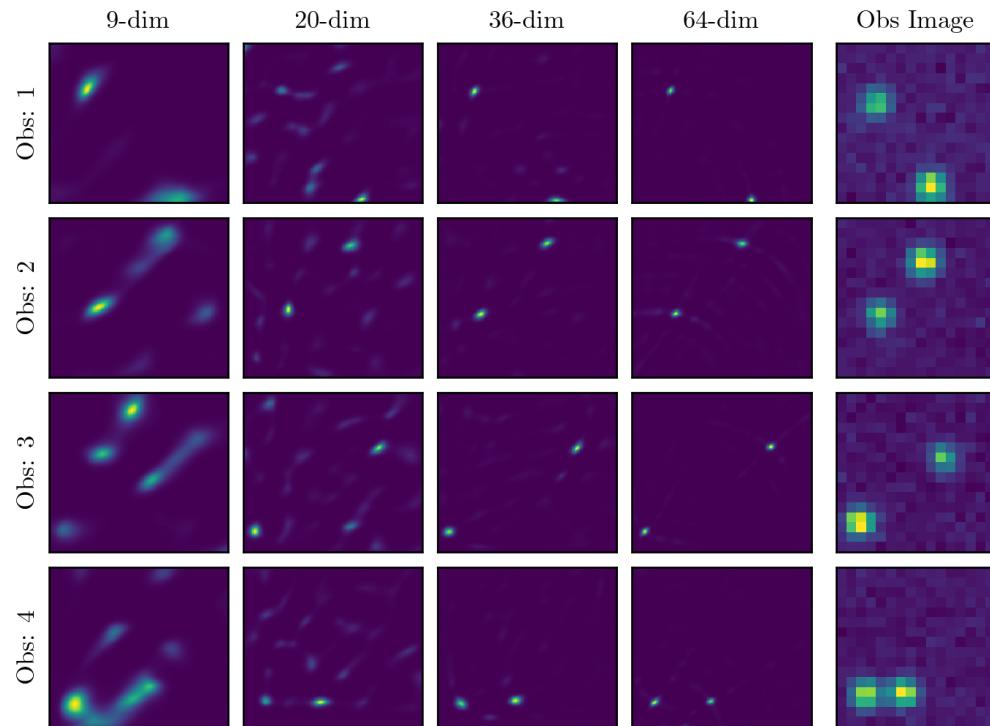
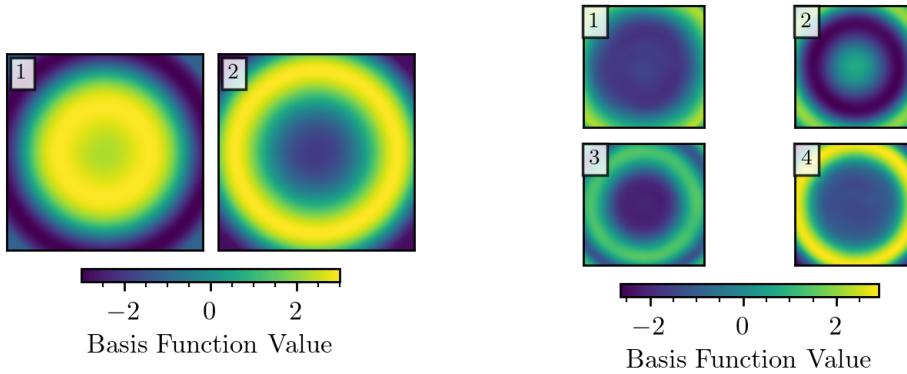
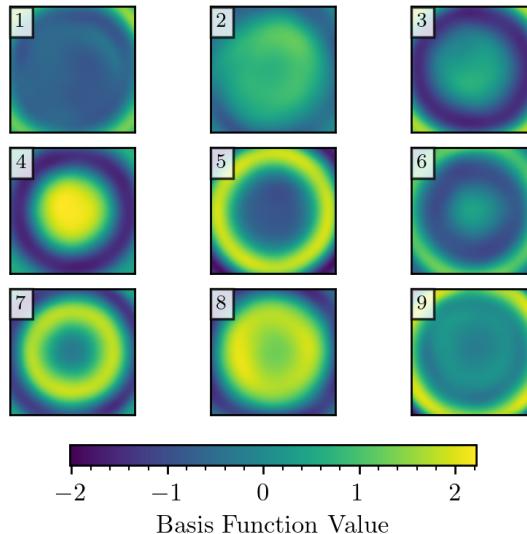


Figure 21: Object detection: 9/20/36/64-dim basis functions and the corresponding estimated posterior density.

1782 Interestingly, in the ring case study, we observe that even a 2-dimensional basis function is sufficient  
 1783 to recover the ring-shaped posterior. As shown in Figure 25, the estimated posterior using a 2-  
 1784 dimensional basis function is visually nearly indistinguishable from the true posterior, with only  
 1785 minor artifacts appearing when the observation  $x$  approaches extreme values (e.g.,  $x = 1.95$ ). This  
 1786 observation is quantitatively supported by the KL divergence results in Table 5, where both forward  
 1787 and reverse KL values are low (0.032/0.031) for the 2-dimensional case. These results demonstrate  
 1788 that for simpler posterior structures, our method can achieve accurate inference with remarkably  
 1789 low-dimensional basis functions.

	2-dim	4-dim	9-dim	20-dim
<b>Forward KL Divergence</b>	0.032	0.0057	0.0056	<b>0.0054</b>
$\Delta$ <b>Forward KL Divergence</b>	-	<b>-0.0263</b>	-0.0001	-0.0002
<b>Reverse KL Divergence</b>	0.031	0.0032	0.0028	<b>0.0027</b>
$\Delta$ <b>Reverse KL Divergence</b>	-	<b>-0.0278</b>	-0.0004	-0.0001

1797 Table 5: Ring: forward/reverse KL divergence for models of different basis function dimensions.  
 17981813 Figure 22: Ring: density plot of 2-dim basis func-  
 1814 tion over plane  $z$ . Each subplot represents the  
 1815 density plot of a certain dimension.Figure 23: Ring: density plot of 4-dim basis func-  
 tion over plane  $z$ .Figure 24: Ring: density plot of 9-dim basis function over plane  $z$ .

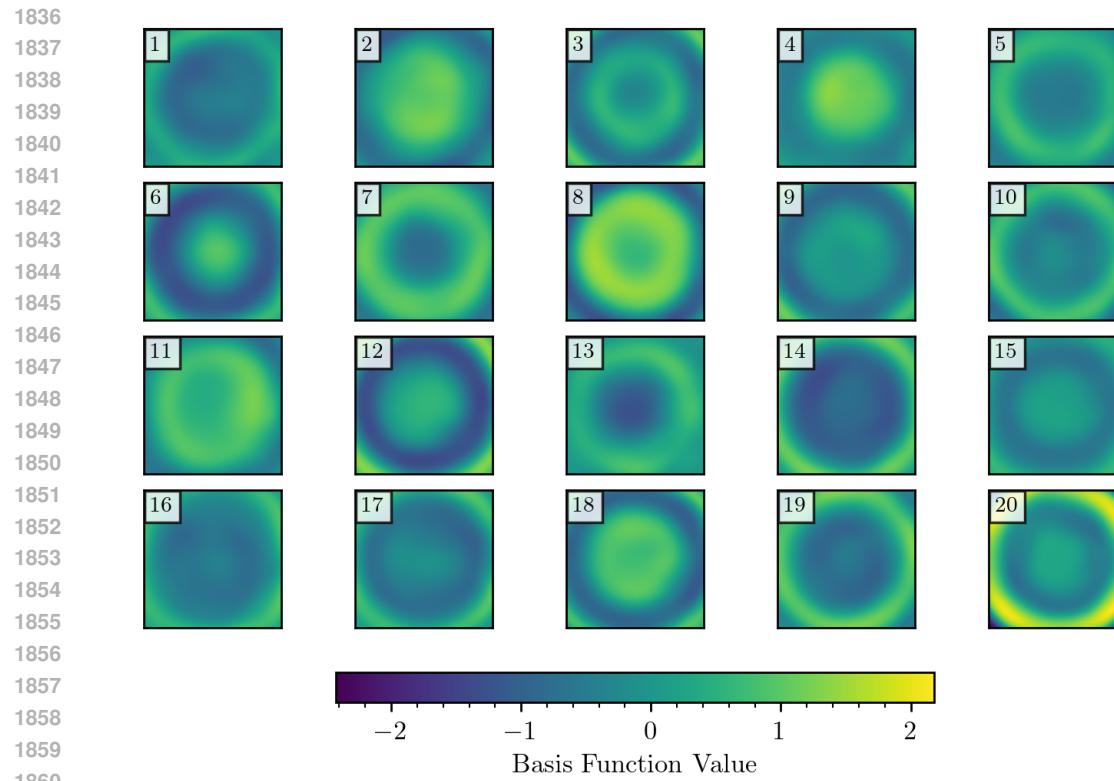
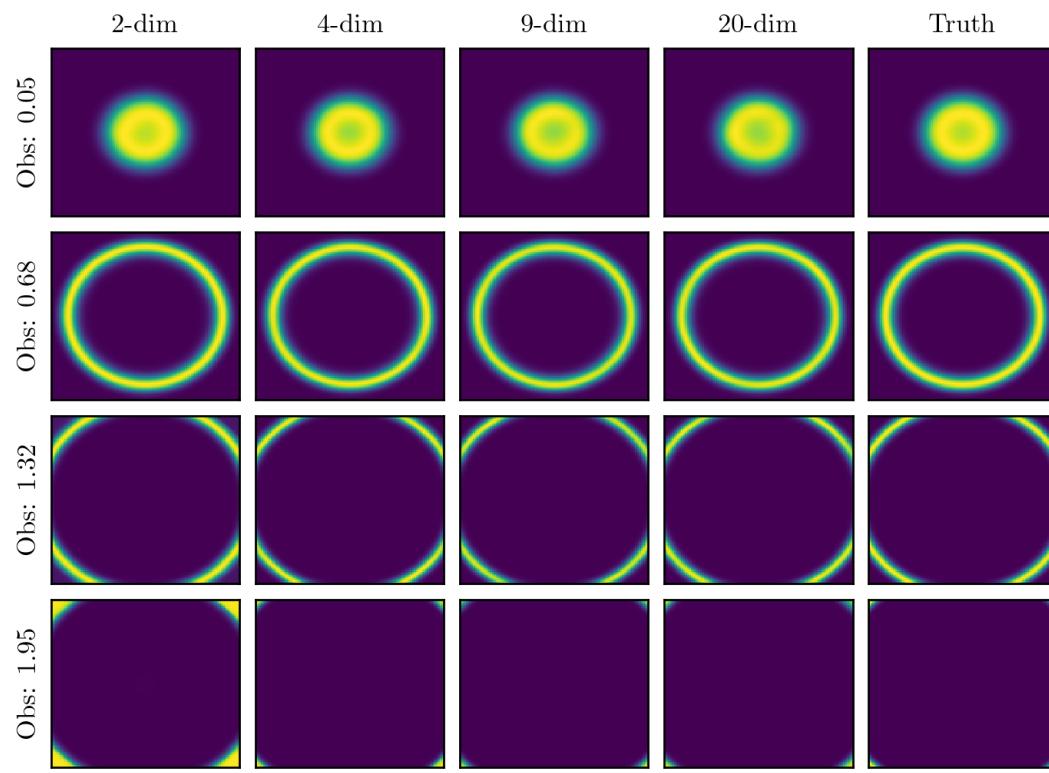
Figure 25: Ring: density plot of 20-dim basis function over plane  $z$ .

Figure 26: Ring: 2/4/9/20-dim basis functions and the corresponding estimated posterior density.

1890

## E.4 TOWARDS HIGH DIMENSION

1891

1892 To evaluate the capability of our model in predicting high-dimensional posteriors, we construct a  
1893 50-dimensional annulus model defined as:

1894

1895

1896 
$$z_1, z_2, \dots, z_{50} \sim \text{Unif}[0, 1],$$
  
1897 
$$z = (z_1, \dots, z_{50}),$$
  
1898 
$$x \mid z \sim \mathcal{N}(\|z\|^2, \sigma^2),$$
  
1899

1900

where  $\sigma^2 = 10^{-2}$ .

1901

1902 In Figures 27 and 28, we randomly select two pairs of dimensions (3, 14) and (43, 47), and visualize  
1903 the estimated posterior over these subspaces, i.e.,  $q(z_3, z_{14} \mid x)$  and  $q(z_{43}, z_{47} \mid x)$ . We discretize  
1904 each pair into a  $100 \times 100$  grid and perform Monte Carlo integration over the remaining 48 dimensions  
1905 to obtain the estimated posterior on the chosen 2D subspace. The results show that the estimated  
1906 posteriors closely match the true posteriors, with minor discrepancies likely due to variance in the  
1907 Monte Carlo integration. The final two columns of each plot display the marginal densities, which  
1908 demonstrate that our model successfully captures the true marginal posterior distributions. This  
1909 good performance can also be verified by quantitative metrics. Our model attains 0.018 forward KL  
1910 divergence and 0.022 reverse KL divergence on average over 50 dimensions.

1911

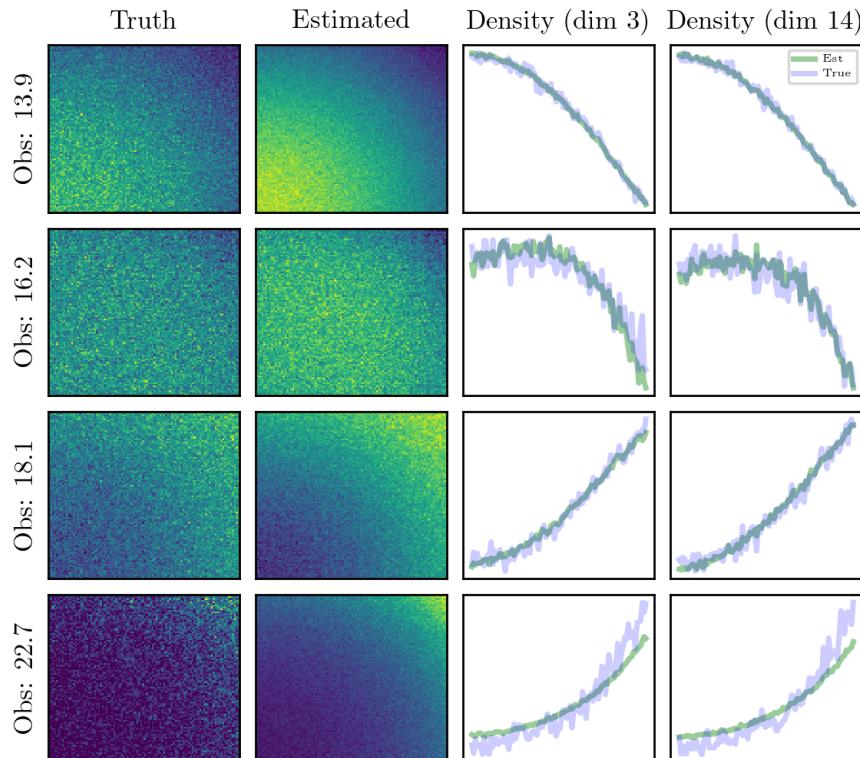


Figure 27: True and estimated posterior density over dimensions 3 and 14. The y-axis in the last two columns represents the marginal density.

1938

1939

1940

1941

1942

1943

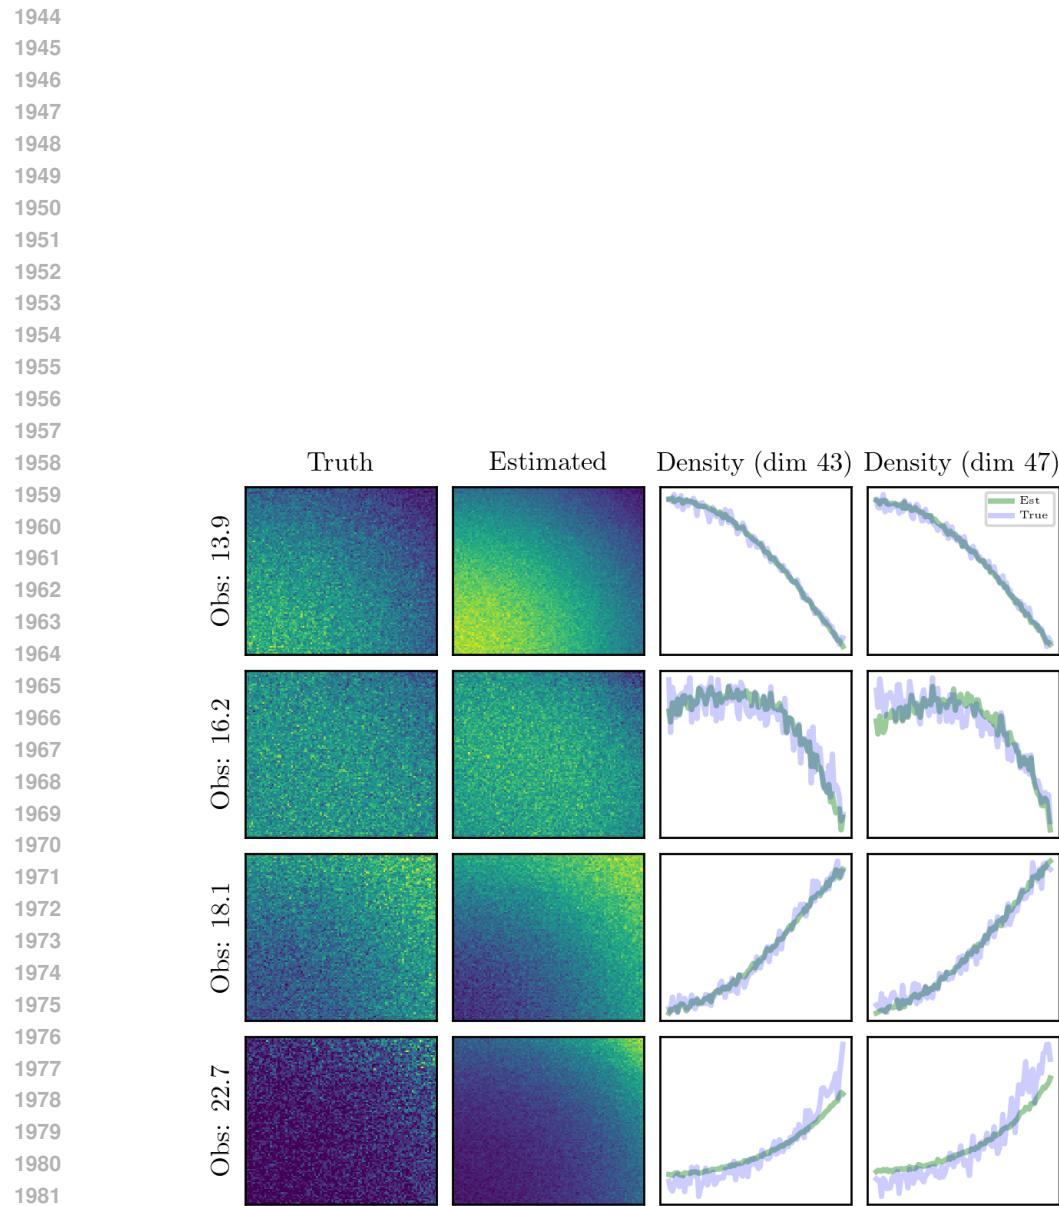


Figure 28: True and estimated posterior density over dimensions 43 and 47. The y-axis in the last two columns represents the marginal density.

1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997

1998  
1999

## E.5 EIGENVI ON 2D CASE STUDIES

2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009

We evaluate EigenVI (Cai et al., 2024) on three two-dimensional targets with thin or curved posterior density patterns. We use a tensor product expansion with  $K = 16$  basis functions per axis ( $K^2 = 256$  coefficients) and 50,000 importance samples for fitting. The reconstructions capture only coarse structure: for the diagonal bands, EigenVI recovers orientation but the two ridges are blurry; for the ring, it fills the central hole and collapses mass inward; for the spiral, it loses the manifold and yields blurry lobes. These failures arise from spectral bias of orthogonal expansions, which under-represent the high-frequency content required by thin or strongly curved posterior. While increasing  $K$  can help, computational and statistical costs scale as  $K^d$  (here  $d = 2$ ), making adequate resolution impractical. In sum, with practical  $K$ , EigenVI is adequate for smooth densities but inadequate for multimodal or topologically nontrivial two-dimensional targets.

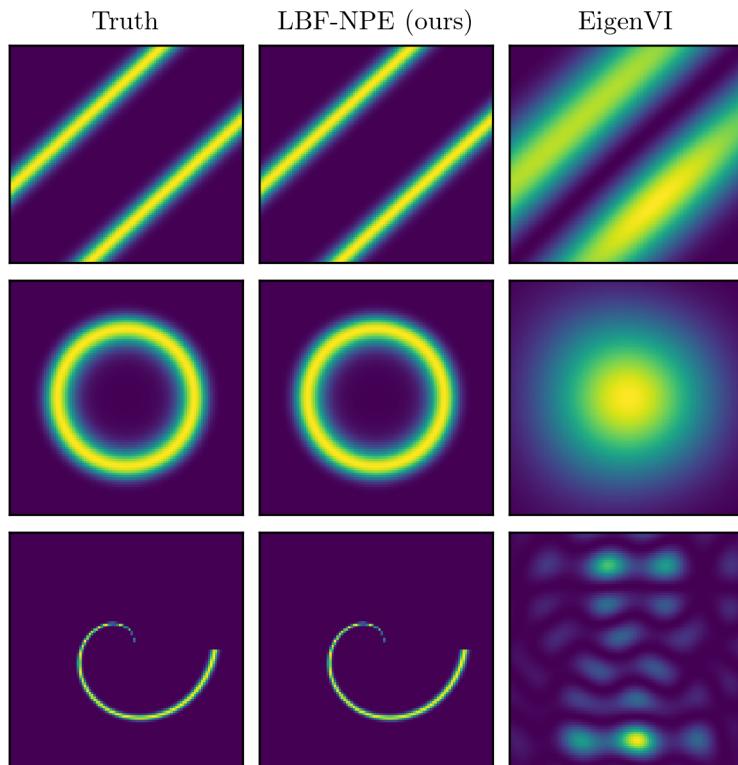
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034

Figure 29: EigenVI results for three 2D test cases. Along each axis, we fit 16 basis functions (i.e.,  $K = 16$  in their original paper); For their importance sampling, we draw 50,000 samples.

2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051

2052  
2053

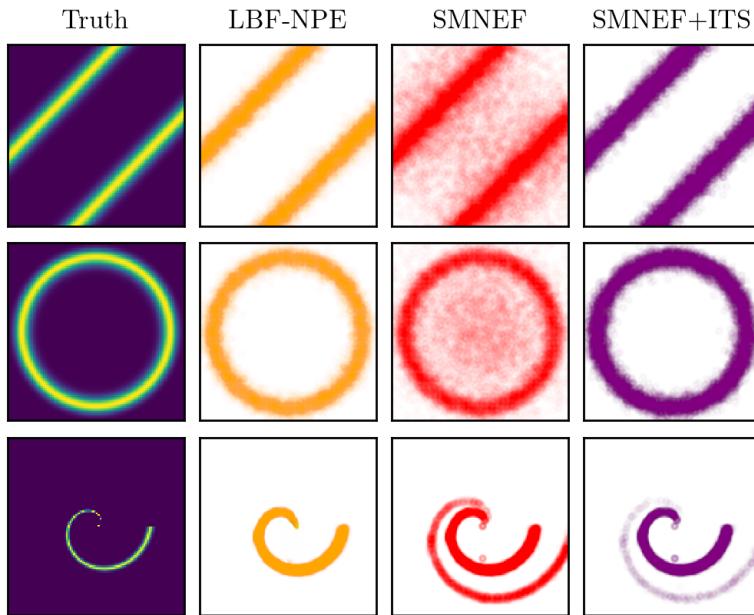
## E.6 SCORE MATCHED NEURAL EXPONENTIAL FAMILIES ON 2D CASE STUDIES

2054

We reproduce the method proposed in Score Matched Neural Exponential Families for Likelihood-Free Inference (Pacchiardi & Dutta, 2022) and evaluate its sampling quality on three two-dimensional case studies. Their approach estimates the unnormalized probability  $p(x | z)$ , in contrast to our focus on  $p(z | x)$ , and employs Exchange MCMC to draw posterior samples. As illustrated in the following figure, the Exchange MCMC samples are suboptimal, often appearing overdispersed and misaligned with the true density. In the bands case study, the samples fail to align with the ridges and instead spread into low-density regions. In the ring case study, many samples are scattered inside the ring rather than concentrating on its boundary where the density peaks. In the spiral case study, the bias is most evident, with samples deviating substantially from the high-density spiral structure. We also apply inverse transform sampling, which was not considered in the original paper. This approach produces samples that more closely follow the true density across all three case studies, though in the spiral example a residual bias is still visible.

2065

2066



2106 E.7 COMPUTATIONAL COST  
2107

2108 Table 6 compares the computational efficiency of several baseline models with the proposed LBF-  
2109 NPE across the 2D case studies, object detection, and redshift estimation. While LBF-NPE shows  
2110 slightly higher per-step runtime and memory usage than some baselines, its principal advantage is  
2111 its substantially faster convergence, requiring markedly fewer training steps, such as 8k in the 2D  
2112 case and 48k in the redshift task. This accelerated convergence results in competitive or superior total  
2113 training time across tasks, indicating that LBF-NPE achieves an effective optimization trajectory. A  
2114 notable detail is that, in the redshift experiments, the computational costs of all three methods appear  
2115 very similar. This is largely due to the dominant overhead from the convolutional U-shape network  
2116 used for image processing, which outweighs differences in the loss computation. Another point of  
2117 clarification is that the GPU memory usage reported here is lower than the values in Section D.2 (e.g.,  
2118 where peak usage is approximately 8400MB). This discrepancy arises because, for the computational  
2119 cost evaluation, we disable GPU memory preallocation in JAX (The JAX Authors, 2024), which  
2120 otherwise reserves roughly 75% of the available GPU memory.

2121 Case study	2122 Method	2123 Time per step (s/step)	2124 GPU Memory (MB)	2125 Converge at (steps)	2126 Time until converge (s)
2123 <b>Sinusoidal likelihood</b> (batch size: 1024)	2124 LBF-NPE	0.038	312	<b>4k</b>	<b>152</b>
	MDN	<b>0.021</b>	<b>248</b>	10k	210
2125 <b>2D case studies</b> (batch size: 1024)	2126 LBF-NPE	0.127	970	<b>8k</b>	1016
	NSF	0.084	780	29k	2436
	RealNVP	0.082	690	25k	2050
	MDN	<b>0.044</b>	<b>430</b>	13k	<b>572</b>
2129 <b>Object detection</b> (batch size: 1024)	2130 LBF-NPE	0.143	2230	15k	2145
2131 <b>Redshift</b> (batch size: 32)	2132 LBF-NPE	0.28	7319	<b>48k</b>	<b>13440</b>
	NSF	0.28	7012	80k	22400
	MDN	<b>0.26</b>	<b>6988</b>	54k	14040

2134 Table 6: Computational cost: For 2D case studies, we only report the computational cost for the  
2135 spiral case study, because the other two case studies have similar computational costs. The  
2136 "Converge at (steps)" refers to the maximum training steps to reach the performance reported in the  
2137 paper. NSF is the abbreviation of Neural Spline Flow.

2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159

2160 E.8 REPEATED EXPERIMENTS  
2161

2162 Across the following four tables, the reported 90% credible intervals (the second term in each table  
2163 entry) demonstrate that our LBF-NPE yields markedly more stable and reliable performance than the  
2164 competing approaches. In the repeated 2D experiments for both forward and reverse KL divergence,  
2165 LBF-NPE consistently attains intervals that are substantially narrower, often by an order of magnitude,  
2166 than those of NSF, RealNVP, and MDN, indicating that our posterior approximations are much less  
2167 sensitive to random initialisation. This advantage persists in the more challenging *spiral* case,  
2168 where LBF-NPE maintains tight credible intervals, while the baselines exhibit markedly inflated  
2169 uncertainty. A similar pattern is observed for the 2D case studies' NLL results. Although NSF  
2170 achieves a slightly better mean NLL on *spiral*, the corresponding credible interval for LBF-NPE  
2171 is smaller, underscoring that our method achieves competitive accuracy with reduced run-to-run  
2172 variability. Finally, in the held-out redshift NLL experiment, the 90% credible interval of LBF-NPE  
2173 is more than a factor of two tighter than those of NSF and MDN, confirming that, beyond achieving  
2174 strong average performance, our method delivers significantly more concentrated and predictable  
2175 outcomes across repeated trials.

	<b>LBF-NPE</b>	<b>NSF</b>	<b>RealNVP</b>	<b>MDN</b>
<b>Bands</b>	<b>0.0048</b> ( $\pm 0.0003$ )	0.016 ( $\pm 0.003$ )	0.015 ( $\pm 0.005$ )	0.182 ( $\pm 0.01$ )
<b>Ring</b>	<b>0.0054</b> ( $\pm 0.0005$ )	0.017 ( $\pm 0.004$ )	0.024 ( $\pm 0.005$ )	0.205 ( $\pm 0.02$ )
<b>Spiral</b>	<b>0.187</b> ( $\pm 0.004$ )	0.201 ( $\pm 0.01$ )	0.545 ( $\pm 0.07$ )	0.948 ( $\pm 0.09$ )

2181 Table 7: Forward KL divergence of LBF-NPE (ours), NSF (Neural Spline Flow), RealNVP, and  
2182 MDN on three 2D test cases. Lower values indicate better posterior approximation.  
2183

	<b>LBF-NPE</b>	<b>NSF</b>	<b>RealNVP</b>	<b>MDN</b>
<b>Bands</b>	<b>0.0014</b> ( $\pm 0.0004$ )	0.0099 ( $\pm 0.001$ )	0.011 ( $\pm 0.007$ )	0.156 ( $\pm 0.02$ )
<b>Ring</b>	<b>0.0027</b> ( $\pm 0.0003$ )	0.013 ( $\pm 0.003$ )	0.014 ( $\pm 0.003$ )	0.204 ( $\pm 0.01$ )
<b>Spiral</b>	<b>0.188</b> ( $\pm 0.005$ )	0.322 ( $\pm 0.04$ )	0.666 ( $\pm 0.09$ )	1.973 ( $\pm 0.14$ )

2184 Table 8: Reverse KL divergence of LBF-NPE (ours), NSF (Neural Spline Flow), RealNVP, and MDN  
2185 on three 2D test cases. Lower values indicate better posterior approximation.  
2186

	<b>LBF-NPE</b>	<b>NSF</b>	<b>RealNVP</b>	<b>MDN</b>
<b>Bands</b>	<b>-0.060</b> ( $\pm 0.07$ )	0.151 ( $\pm 0.23$ )	0.157 ( $\pm 0.22$ )	1.389 ( $\pm 0.41$ )
<b>Ring</b>	<b>0.030</b> ( $\pm 0.03$ )	0.621 ( $\pm 0.24$ )	0.733 ( $\pm 0.11$ )	1.031 ( $\pm 0.18$ )
<b>Spiral</b>	0.838 ( $\pm 0.13$ )	<b>0.727</b> ( $\pm 0.25$ )	0.859 ( $\pm 0.32$ )	2.788 ( $\pm 0.31$ )

2187 Table 9: Negative log-likelihood (NLL) of LBF-NPE (ours), NSF (Neural Spline Flow), RealNVP,  
2188 and MDN on three 2D test cases. Lower values indicate better posterior approximation.  
2189

	<b>LBF-NPE</b>	<b>NSF</b>	<b>MDN</b>
<b>NLL</b>	<b>-57,220</b> ( $\pm 152$ )	-55,389 ( $\pm 379$ )	-50,648 ( $\pm 322$ )

2200 Table 10: Held-out NLL of the true redshift  $z$ . NSF is the abbreviation of Neural Spline Flow.  
2201

2214  
2215

## E.9 BAYESIAN NEURAL NETWORKS AS GENERATIVE MODELS

2216  
2217  
2218  
2219  
2220  
2221

A Bayesian neural network (BNN) is a neural network in which each weight (and bias) is treated as a probability distribution rather than a fixed value (MacKay, 1992a;b; Neal, 1996). When making predictions, it marginalizes over these distributions to produce not just a prediction, but also an estimate of uncertainty. We consider BNNs as generative models to illustrate that LBF-NPE can perform posterior predictive inference while implicitly marginalizing over a high-dimensional parameter space. In particular, we consider a two-layer fully connected BNN:

2222  
2223

$$\text{BNN}_\theta(x) = \text{Linear}_{\theta_2}(\text{ReLU}(\text{Linear}_{\theta_1}(x))), \quad \theta_1, \theta_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (22)$$

2224  
2225  
2226  
2227

where  $\theta_1$  and  $\theta_2$  are the weight matrices of the first and second linear layers, respectively, and  $\text{ReLU}(\cdot)$  denotes the rectified linear unit activation. The network takes a one-dimensional input and produces a one-dimensional output through a hidden layer of width 16. Throughout this experiment we restrict attention to  $x \in [0, 10]$  and  $y \in [-8, 8]$ .

2228

Let  $\theta = (\theta_1, \theta_2)$  and let the dataset be

2229  
2230

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n, \quad y_i = \text{BNN}_\theta(x_i) + \epsilon_i, \quad n = 5,$$

2231  
2232

for a fixed draw of  $\theta$  from  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , and noise  $\epsilon_i \sim \mathcal{N}(0, 1)$ . The posterior predictive distribution for a new pair  $(x', y')$  is

2233  
2234

$$p(y' | x', \mathcal{D}) = \int p(y' | x', \theta) p(\theta | \mathcal{D}) d\theta, \quad (23)$$

2235  
2236  
2237

which involves integration over the high-dimensional weight vector  $\theta$ . Our goal in this section is to show that LBF-NPE can approximate the conditional density in (23) without ever explicitly sampling or optimizing over  $\theta$ .

2238  
2239  
2240  
2241

In our implementation, LBF-NPE parameterizes the conditional density via a basis-function network  $s_\psi$  and a coefficient network  $f_\phi$ . The coefficient network  $f_\phi$  takes as input the query variate  $x'$  together with the conditioning set  $\mathcal{D}$ , while the basis-function network  $s_\psi$  takes  $y'$  as input. Together they define

2242

$$\text{LBF-NPE}(y', x', \mathcal{D}) \approx p(y' | x', \mathcal{D}),$$

2243  
2244  
2245  
2246

and are trained using Algorithm 1. Both  $s_\psi$  and  $f_\phi$  are implemented as four-layer multilayer perceptrons with 128 hidden units per layer; the resulting basis-function and coefficient vectors have dimension 20. We optimize the parameters  $(\psi, \phi)$  with the AdamW optimizer (Loshchilov & Hutter, 2019), using a learning rate of  $10^{-3}$  for 10,000 gradient steps.

2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255

The qualitative behavior of the learned posterior predictive distributions is shown in Figure 31. Across the four panels, the underlying BNN functions (orange curves) exhibit markedly different slopes, curvatures, and ranges, yet LBF-NPE recovers their overall shape from only  $n = 5$  observations. Moreover, the posterior highest-density interval (HDI) is wide in regions with sparse observations and narrow in regions with many data points. For example, when no observation is available near  $x \in [0, 2]$  (top-left panel), the HDI is wide, whereas in regions densely populated with observations (e.g.,  $x \in [2.5, 5.0]$  in the bottom-left panel) the HDI becomes narrow. This pattern indicates that LBF-NPE successfully captures the epistemic uncertainty induced by marginalization over the BNN weights.

2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267

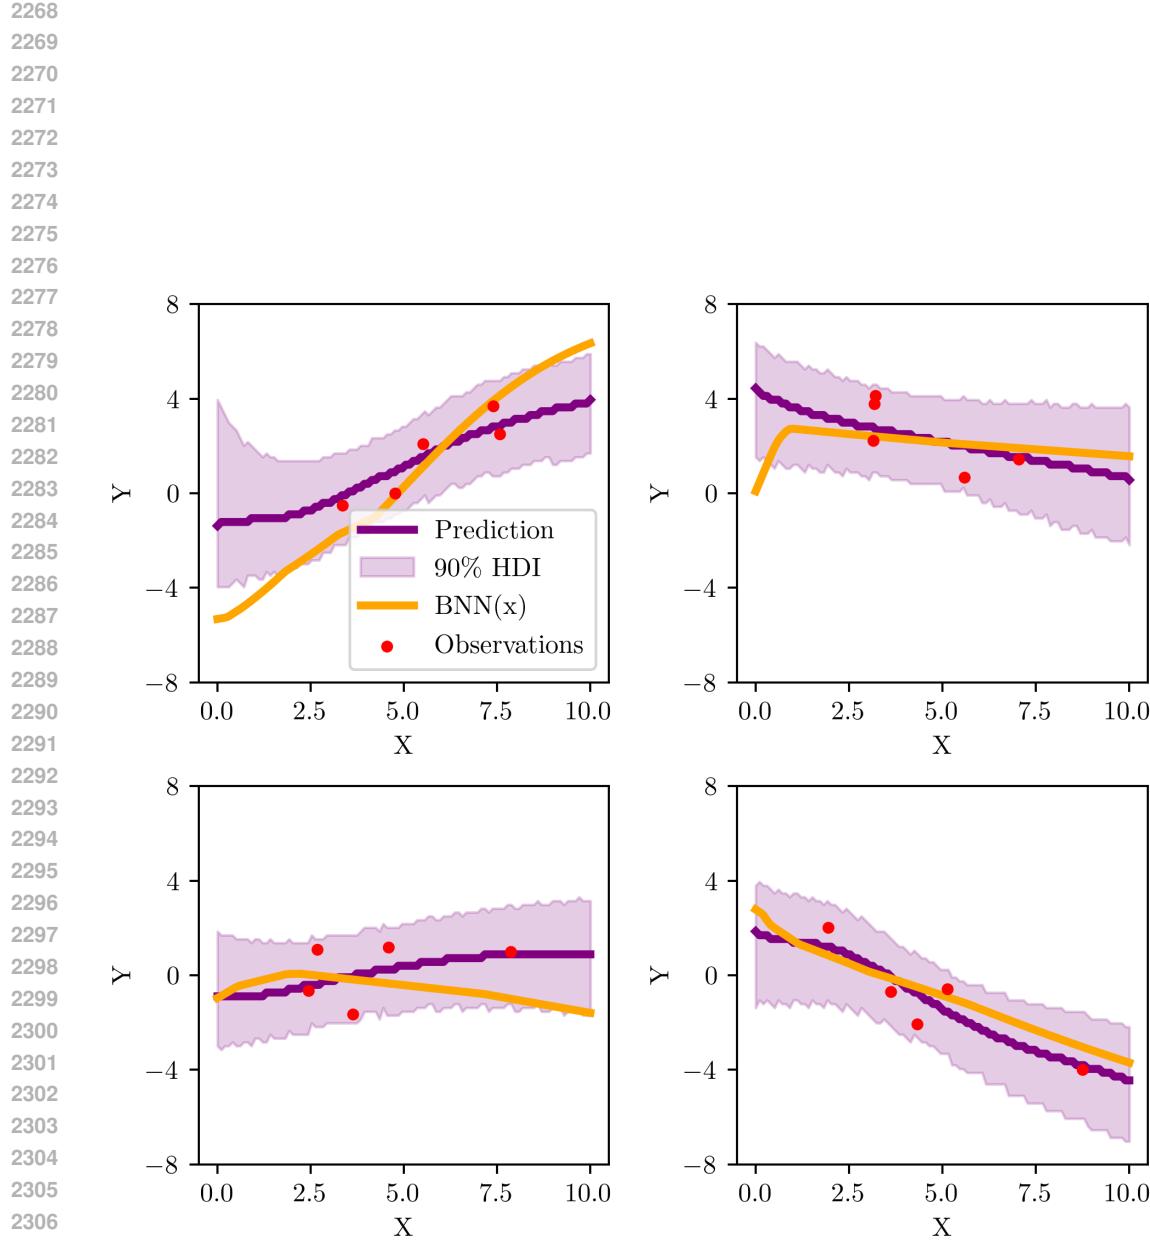


Figure 31: Posterior predictive distributions obtained from LBF-NPE on four synthetic regression problems generated by a Bayesian neural network (BNN). Each panel corresponds to a different draw of the BNN weights. The orange curve is the true mapping  $x \mapsto BNN_\theta(x)$ , and the red dots denote the  $n = 5$  observed data points used for inference. The solid purple curve shows the pointwise mode of the posterior predictive density produced by LBF-NPE, and the shaded region indicates the associated 90% highest-density interval (HDI).