

---

# On Contrastive Learning for Likelihood-free Inference

---

Conor Durkan<sup>1</sup> Iain Murray<sup>1</sup> George Papamakarios<sup>2</sup>

## Abstract

Likelihood-free methods perform parameter inference in stochastic simulator models where evaluating the likelihood is intractable but sampling synthetic data is possible. One class of methods for this likelihood-free problem uses a classifier to distinguish between pairs of parameter-observation samples generated using the simulator and pairs sampled from some reference distribution, which implicitly learns a density ratio proportional to the likelihood. Another popular class of methods fits a conditional distribution to the parameter posterior directly, and a particular recent variant allows for the use of flexible neural density estimators for this task. In this work, we show that both of these approaches can be unified under a general contrastive learning scheme, and clarify how they should be run and compared.

## 1. Introduction

Modeling systems using parameterized stochastic simulators is prevalent across many scientific and engineering disciplines, including cosmology (Alsing et al., 2019), high-energy physics (Brehmer et al., 2018a), and computational neuroscience (Gonçalves et al., 2019). Specifying models in this way is appealing since it is often easier to describe a generative process implicitly rather than reasoning directly about an emergent probability distribution. Traditional parameter inference algorithms such as variational methods and Markov Chain Monte Carlo (MCMC) usually don't apply to these models, since explicit evaluation of the likelihood function is often intractable.

In recent years, significant progress has been made toward this challenge of likelihood-free inference (Sisson et al., 2018). Most of the literature relies on sample-based methods that require very large numbers of simulations. Methods

---

<sup>1</sup>School of Informatics, University of Edinburgh, United Kingdom <sup>2</sup>DeepMind, London, United Kingdom. Correspondence to: Conor Durkan <conor.durkan@ed.ac.uk>.

based on neural networks have reduced the number of simulations needed to perform effective inference by orders of magnitude (Papamakarios, 2019; Cranmer et al., 2019). Some of these can evaluate and sample from an approximate posterior directly, bypassing the need for MCMC schemes altogether. Ultimately, this line of work seeks to develop a robust and practical toolkit which practitioners can rely upon for whichever use case they desire.

In this paper, we consider two recently proposed approaches to neural likelihood-free inference: the first uses classification to approximate density ratios proportional to the likelihood (Hermans et al., 2020), whereas the second directly casts the problem as a conditional density estimation task (Greenberg et al., 2019). We demonstrate that these two methods, traditionally viewed as distinct and compared as such in the literature, are both instances of a more general contrastive learning scheme, and can thus be unified under a single framework. Using this perspective, we directly compare the properties and behaviour of both algorithms, and offer practical recommendations for those wishing to use these methods.

## 2. Background

Given a vector of parameters  $\theta$ , a stochastic simulator generates latent random numbers  $\mathbf{z}$ , and produces observed data  $\mathbf{x} = \mathbf{g}(\theta, \mathbf{z})$ , where  $\mathbf{g}$  is some nonlinear function. The likelihood of the parameters  $\theta$  given observed data  $\mathbf{x}$  is then

$$p(\mathbf{x} | \theta) = \int \delta(\mathbf{x} - \mathbf{g}(\theta, \mathbf{z})) p(\mathbf{z} | \theta) d\mathbf{z}, \quad (1)$$

where  $\delta(\cdot)$  is the Dirac delta, and this integral is intractable in general. In some cases it may be possible to evaluate the joint likelihood  $p(\mathbf{x}, \mathbf{z} | \theta)$ , but (i) this requires performing inference on the constrained manifold  $(\mathbf{g}(\theta, \mathbf{z}), \mathbf{z})$ , and (ii) often this augmented space is sufficiently high-dimensional to make standard inference approaches infeasible. In other cases, the simulator may be provided to us as a black-box whose internal workings are not accessible, and it is still desirable to carry out inference in this scenario.

We are interested in computing the posterior  $p(\theta | \mathbf{x}_0)$  for a specific observation  $\mathbf{x}_0$ , where we have prior beliefs  $p(\theta)$  about the parameters. This amounts to computing the density ratio  $\frac{p(\mathbf{x}_0 | \theta)}{p(\mathbf{x}_0)} = \frac{p(\theta | \mathbf{x}_0)}{p(\theta)}$ , and this quantity need only be

known up to a constant independent of  $\theta$  if our approach uses MCMC schemes. Throughout, we assume that the simulator defines a valid probability density function  $p(\mathbf{x} | \theta)$  over observations  $\mathbf{x}$ .

## 2.1. Learning density ratios

It is well known that binary classification can be used to learn density ratios (Hastie et al., 2001; Sugiyama et al., 2012; Mohamed & Lakshminarayanan, 2016). Defining  $p_{\text{pos}}(\mathbf{x}) = p(\mathbf{x} | y=1)$  and  $p_{\text{neg}}(\mathbf{x}) = p(\mathbf{x} | y=0)$  for binary  $y$ , and assuming each class is equally likely a priori yields

$$p(y = 1 | \mathbf{x}) = \frac{p_{\text{pos}}(\mathbf{x})}{p_{\text{pos}}(\mathbf{x}) + p_{\text{neg}}(\mathbf{x})} = \frac{r(\mathbf{x})}{r(\mathbf{x}) + 1}, \quad (2)$$

where  $r(\mathbf{x}) = \frac{p_{\text{pos}}(\mathbf{x})}{p_{\text{neg}}(\mathbf{x})}$ . In other words, an optimal binary classifier which distinguishes between samples from the two distributions recovers the ratio between their density functions. A common choice of parameterization for this classifier is  $p(y = 1 | \mathbf{x}) = \sigma(f_\phi(\mathbf{x}))$ , where  $\sigma(\cdot)$  is the logistic sigmoid and  $f_\phi$  is a real-valued function whose parameters  $\phi$  can be fit using maximum likelihood, and the optimal value for  $f_\phi$  is  $f_\phi(\mathbf{x}) = \log r(\mathbf{x})$ .

In the context of likelihood-free inference, Hermans et al. (2020) consider positive examples ( $y = 1$ ) those pairs  $(\mathbf{x}, \theta) \sim p(\mathbf{x}, \theta)$  which are distributed jointly, and generated by sampling a parameter from the prior and subsequently simulating an observation for that parameter setting. In contrast, negative examples ( $y = 0$ ) are those pairs  $(\mathbf{x}, \theta) \sim p(\mathbf{x}) p(\theta)$  with parameters and observations sampled independently from their respective marginal distributions. Then binary classification recovers the ratio

$$r(\mathbf{x}, \theta) = \frac{p(\mathbf{x}, \theta)}{p(\mathbf{x}) p(\theta)} = \frac{p(\mathbf{x} | \theta)}{p(\mathbf{x})} = \frac{p(\theta | \mathbf{x})}{p(\theta)}. \quad (3)$$

Since the prior is known, it is possible to evaluate the parameter posterior exactly using  $p(\theta | \mathbf{x}_0) = r(\mathbf{x}_0, \theta) p(\theta)$ . However, only exact evaluation of the posterior distribution is achieved in this case, whereas sampling still requires an MCMC scheme. Likelihood-free inference by ratio estimation (LFIRE, Thomas et al., 2016) is a pioneering example of this approach. However, Thomas et al. (2016) fit a classifier using only observations  $\mathbf{x}$  as input, meaning a separate classifier is required for each posterior evaluation  $p(\theta | \mathbf{x})$ .

Moreover, Hermans et al. (2020) note that since we are interested in computing the posterior for a particular observation  $\mathbf{x}_0$ , it may be wasteful to learn the ratio  $r(\mathbf{x}, \theta)$  for all pairs  $(\mathbf{x}, \theta)$ . Instead, it might be preferable to focus attention on those parameters which may have plausibly generated  $\mathbf{x}_0$ , rather than parameters sampled from the prior. Suppose we have access to such a proposal prior, which we denote  $\tilde{p}(\theta)$ . Then, considering positive examples

$(\mathbf{x}, \theta) \sim \tilde{p}(\mathbf{x}, \theta) = p(\mathbf{x} | \theta) \tilde{p}(\theta)$  and negative examples  $(\mathbf{x}, \theta) \sim \tilde{p}(\mathbf{x}) \tilde{p}(\theta)$ , binary classification recovers the ratio

$$\tilde{r}(\mathbf{x}, \theta) = \frac{\tilde{p}(\mathbf{x}, \theta)}{\tilde{p}(\mathbf{x}) \tilde{p}(\theta)} = \frac{p(\mathbf{x} | \theta)}{\tilde{p}(\mathbf{x})} \quad (4)$$

At the expense of exact posterior evaluation, we have ostensibly gained a more accurate ratio estimate for the parameters corresponding to the observation of interest. Moreover, it is still possible to evaluate the posterior up to a constant, since  $p(\theta | \mathbf{x}_0) \propto \tilde{r}(\mathbf{x}_0, \theta) p(\theta)$ , so MCMC schemes can still be used to generate samples from the desired posterior. Using this approach, an algorithm emerges where the likelihood ratio is refined iteratively across a series of rounds, and training data for a given round is generated using the estimated ratios from preceding rounds i.e. the proposal  $\tilde{p}(\theta)$  is a mixture of posterior estimates from all preceding rounds. We use the term Sequential Ratio Estimation (SRE) to describe this approach of Hermans et al. (2020).

Finally, binary classification can equivalently be viewed as a task where the model must correctly identify the correct example given a set consisting of exactly one positive example drawn jointly with an observation, and one negative example drawn marginally. This equivalent perspective will prove useful in deriving a unifying notation in later sections. Denoting this set  $\Theta = \{\theta_0, \theta_1\}$  and assuming each class is equally likely a priori, we have

$$p(y = 1 | \mathbf{x}, \Theta) = \frac{p(\mathbf{x}, \Theta | y = 1)}{p(\mathbf{x}, \Theta | y = 0) + p(\mathbf{x}, \Theta | y = 1)} \quad (5)$$

$$= \frac{p(\theta_1 | \mathbf{x}) p(\theta_0)}{p(\theta_0 | \mathbf{x}) p(\theta_1) + p(\theta_1 | \mathbf{x}) p(\theta_0)} \quad (6)$$

$$= \frac{p(\theta_1 | \mathbf{x}) / p(\theta_1)}{p(\theta_0 | \mathbf{x}) / p(\theta_0) + p(\theta_1 | \mathbf{x}) / p(\theta_1)}. \quad (7)$$

Now if we parameterize a binary classifier by

$$p(y = 1 | \mathbf{x}, \Theta) = \frac{\exp(f_\phi(\theta_1, \mathbf{x}))}{\exp(f_\phi(\theta_0, \mathbf{x})) + \exp(f_\phi(\theta_1, \mathbf{x}))}, \quad (8)$$

then the optimal classifier has  $f_\phi(\theta, \mathbf{x}) = \log \frac{p(\theta | \mathbf{x})}{p(\theta)} + c(\mathbf{x})$ , where  $c(\mathbf{x})$  is constant with respect to  $\theta$ .

## 2.2. Learning the posterior

Likelihood-free inference can also be cast as a conditional density estimation problem directly. Consider a parameterized conditional density estimator  $q_\phi(\theta | \mathbf{x})$ , sometimes referred to as a recognition model. The parameters  $\phi$  can be fit by minimizing  $\mathcal{L}(\phi) = \mathbb{E}_{(\mathbf{x}, \theta) \sim p(\mathbf{x}, \theta)} [-\log q_\phi(\theta | \mathbf{x})]$ . With enough data and a sufficiently flexible model,  $q_\phi(\theta | \mathbf{x})$  will converge to the posterior density  $p(\theta | \mathbf{x})$  for all pairs  $(\mathbf{x}, \theta)$ . Then the posterior density of interest can be computed using  $p(\theta | \mathbf{x}_0) = q_\phi(\theta | \mathbf{x}_0)$ . However, since we are

again only interested in the posterior for a specific observation  $\mathbf{x}_0$ , it may be beneficial to learn the conditional density more carefully for parameters which could have plausibly generated that observation. Indeed, recent work, which we group under the umbrella of Sequential Neural Posterior Estimation (SNPE, Papamakarios & Murray, 2016; Lueckmann et al., 2017; Greenberg et al., 2019), has developed sequential procedures for exactly this purpose, narrowing focus to parameters of interest across a series of rounds, where the posterior estimate for a given round is used as the proposal prior for the subsequent round. Central to the function of these methods is the proposal posterior  $\tilde{p}(\boldsymbol{\theta} | \mathbf{x})$ , related to the true posterior by

$$\tilde{p}(\boldsymbol{\theta} | \mathbf{x}) = p(\boldsymbol{\theta} | \mathbf{x}) \frac{\tilde{p}(\boldsymbol{\theta}) p(\mathbf{x})}{p(\boldsymbol{\theta}) \tilde{p}(\mathbf{x})}. \quad (9)$$

**SNPE-A.** The first method is known as SNPE-A (Papamakarios & Murray, 2016). Given a proposal prior  $\tilde{p}(\boldsymbol{\theta})$ , SNPE-A fits a neural density estimator  $q_\phi(\boldsymbol{\theta} | \mathbf{x})$  by minimizing  $\mathcal{L}(\boldsymbol{\theta}) = \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x} | \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\theta})} [-\log q_\phi(\boldsymbol{\theta} | \mathbf{x})]$ , converging to the proposal posterior  $\tilde{p}(\boldsymbol{\theta} | \mathbf{x})$ . The relation  $p(\boldsymbol{\theta} | \mathbf{x}) \propto \tilde{p}(\boldsymbol{\theta} | \mathbf{x}) \frac{p(\boldsymbol{\theta})}{\tilde{p}(\boldsymbol{\theta})} = q_\phi(\boldsymbol{\theta} | \mathbf{x}) \frac{p(\boldsymbol{\theta})}{\tilde{p}(\boldsymbol{\theta})}$  can then be used to compute the desired conditional density up to a constant. In certain cases, we can compute the posterior analytically. For example, when both  $p(\boldsymbol{\theta})$  and  $\tilde{p}(\boldsymbol{\theta})$  are Gaussian, and the density estimator  $q_\phi(\boldsymbol{\theta} | \mathbf{x})$  is a mixture density network (Bishop, 1994) with  $M$  Gaussian components, the posterior is also a mixture of Gaussians with  $M$  components. However, the algorithm precludes the use of more flexible conditional density estimators, such as autoregressive models (Uribe et al., 2016) or normalizing flows (Papamakarios et al., 2019a), if we want a tractable posterior.

**SNPE-B.** The second of these methods, referred to as SNPE-B (Lueckmann et al., 2017), allows for the use of arbitrary neural density estimators. This is achieved through minimizing an importance-weighted loss  $\mathcal{L}(\phi) = \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x} | \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\theta})} \left[ -\frac{p(\boldsymbol{\theta})}{\tilde{p}(\boldsymbol{\theta})} \log q_\phi(\boldsymbol{\theta} | \mathbf{x}) \right]$ , where  $q_\phi(\boldsymbol{\theta} | \mathbf{x})$  now recovers the correct posterior directly. Unfortunately, the importance weights  $\frac{p(\boldsymbol{\theta})}{\tilde{p}(\boldsymbol{\theta})}$  are often high variance, and can lead to poor performance of the algorithm overall.

**SNPE-C (APT).** Finally, the third method, and of particular interest in this paper, is Automatic Posterior Transformation (APT, Greenberg et al., 2019), or SNPE-C in the existing taxonomy. SNPE-C neatly reparameterizes the proposal posterior objective from SNPE-A so that maximum likelihood recovers the true posterior directly. Setting

$$\tilde{q}_\phi(\boldsymbol{\theta} | \mathbf{x}) = q_\phi(\boldsymbol{\theta} | \mathbf{x}) \frac{\tilde{p}(\boldsymbol{\theta})}{p(\boldsymbol{\theta})} \frac{1}{Z_\phi(\mathbf{x})}, \quad (10)$$

where  $Z_\phi(\mathbf{x}) = \int q_\phi(\boldsymbol{\theta} | \mathbf{x}) \frac{\tilde{p}(\boldsymbol{\theta})}{p(\boldsymbol{\theta})} d\boldsymbol{\theta}$ , SNPE-C minimizes  $\mathcal{L}(\phi) = \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x} | \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\theta})} [-\log \tilde{q}_\phi(\boldsymbol{\theta} | \mathbf{x})]$ .

Unfortunately, SNPE-C as described requires the ability to calculate the normalizing constant  $Z_\phi(\mathbf{x})$ , which can be computed analytically for an MDN, but not in general for more flexible density estimators such as normalizing flows. However, Greenberg et al. (2019) note that if the proposal is uniform and supported on a finite set of discrete ‘atoms’  $\Theta$ , the integral reduces to a sum, and the proposal posterior now defines a distribution over this discrete set:

$$\tilde{q}_\phi(\boldsymbol{\theta} | \mathbf{x}) = \frac{q_\phi(\boldsymbol{\theta} | \mathbf{x})/p(\boldsymbol{\theta})}{\sum_{\boldsymbol{\theta}' \in \Theta} q_\phi(\boldsymbol{\theta}' | \mathbf{x})/p(\boldsymbol{\theta}')}. \quad (11)$$

Moreover, Greenberg et al. (2019) show that when the contrasting set is drawn from a distribution over parameters whose support contains the support of the true posterior, the density estimator  $q_\phi(\boldsymbol{\theta} | \mathbf{x})$  still converges to the full continuous posterior. Intuitively, the inference problem has been rephrased as a series of multiple choice questions: given an observation  $\mathbf{x}$  and a set of possible parameters  $\Theta$  which may have generated that observation, the task is to identify the correct parameter  $\boldsymbol{\theta}$ . For the rest of the paper, SNPE-C refers specifically to this atomic variant of the algorithm.

### 3. Method

**Unifying SRE and SNPE-C.** There is clear similarity between eq. (8) and eq. (11); indeed, modulo the switch from a binary to multi-class setting, the difference is only notational. SRE and SNPE-C are in fact both instances of a more general contrastive learning scheme. To see this, consider the posterior over labels for  $K$  categories given an observation  $\mathbf{x}$  and a set of examples  $\Theta$  containing exactly one positive example  $\boldsymbol{\theta}$  which generated that observation, while the remaining set of contrasting examples  $\Theta^{(\setminus k)}$  are drawn jointly from some other distribution conditioned on this joint pair. Then, again assuming each class is equally likely a priori, we have

$$\begin{aligned} p(y = k | \mathbf{x}, \Theta) &= \frac{p(\mathbf{x}, \Theta | y = k) p(y = k)}{\sum_{k'} p(\mathbf{x}, \Theta | y = k') p(y = k')} \quad (12) \\ &= \frac{p(\boldsymbol{\theta}^{(k)} | \mathbf{x}) p(\Theta^{(\setminus k)} | \mathbf{x}, \boldsymbol{\theta}^{(k)})}{\sum_{k'} p(\boldsymbol{\theta}^{(k')} | \mathbf{x}) p(\Theta^{(\setminus k')} | \mathbf{x}, \boldsymbol{\theta}^{(k')})}. \end{aligned}$$

Now suppose the elements of the contrasting set are drawn independently of both the joint pair  $(\mathbf{x}, \boldsymbol{\theta}^{(k)})$ , and of each other, by sampling  $K - 1$  times from the prior. In this case,  $p(\Theta^{(\setminus k)} | \mathbf{x}, \boldsymbol{\theta}^{(k)}) = \prod_{l \neq k} p(\boldsymbol{\theta}^{(l)})$ , and we have

$$p(y = k | \mathbf{x}, \Theta) = \frac{p(\boldsymbol{\theta}^{(k)} | \mathbf{x}) \prod_{l \neq k} p(\boldsymbol{\theta}^{(l)})}{\sum_{k'} p(\boldsymbol{\theta}^{(k')} | \mathbf{x}) \prod_{l \neq k'} p(\boldsymbol{\theta}^{(l)})} \quad (13)$$

$$= \frac{p(\boldsymbol{\theta}^{(k)} | \mathbf{x})/p(\boldsymbol{\theta}^{(k)})}{\sum_{k'} p(\boldsymbol{\theta}^{(k')} | \mathbf{x})/p(\boldsymbol{\theta}^{(k')})} \quad (14)$$

$$= \frac{\exp(f_\phi(\boldsymbol{\theta}^{(k)}, \mathbf{x}) + c(\mathbf{x}))}{\sum_{k'} \exp(f_\phi(\boldsymbol{\theta}^{(k')}, \mathbf{x}) + c(\mathbf{x}))}, \quad (15)$$

---

**Algorithm 1** (Sequential) Contrastive Likelihood-free Inference

---

- 1: **Input:** Simulator  $p(\mathbf{x}|\boldsymbol{\theta})$ , prior  $p(\boldsymbol{\theta})$ , true observation  $\mathbf{x}_0$ , classifier  $f_\phi(\boldsymbol{\theta}, \mathbf{x})$ , rounds  $R$ , simulations per round  $N$ .  
 {For SNPE-C, we require  $f_\phi(\boldsymbol{\theta}, \mathbf{x}) = \log \frac{q_\phi(\boldsymbol{\theta}|\mathbf{x})}{p(\boldsymbol{\theta})}$ , where  $q_\phi(\boldsymbol{\theta}|\mathbf{x})$  is a normalized conditional density estimator.}
  - 2: **Initialize:** Posterior  $p^{(1)}(\boldsymbol{\theta}) = p(\boldsymbol{\theta})$ , dataset  $\mathcal{D} = \{\}$ .
  - 3: **for**  $r = 1$  **to**  $R$  **do**
  - 4:   Sample  $\boldsymbol{\theta}^{(n)} \sim p^{(r)}(\boldsymbol{\theta})$ ,  $n = 1 \dots N$ . {For SRE, this step requires MCMC, but in the case of SNPE-C we can sample i.i.d. from the density estimator, using rejection sampling if the density estimator has support outside the prior.}
  - 5:   Simulate  $\mathbf{x}^{(n)} \sim p(\mathbf{x}|\boldsymbol{\theta}^{(n)})$ ,  $n = 1 \dots N$ .
  - 6:   Aggregate training data  $\mathcal{D} = \mathcal{D} \cup \{(\mathbf{x}^{(n)}, \boldsymbol{\theta}^{(n)})\}_{n=1}^N$ .
  - 7:   **while**  $f_\phi$  not converged **do**
  - 8:     Sample mini-batch  $\{(\mathbf{x}^{(b)}, \boldsymbol{\theta}^{(b)})\}_{b=1}^B \sim \mathcal{D}$ . {This is equivalent to sampling parameters from a proposal  $\boldsymbol{\theta} \sim \tilde{p}(\boldsymbol{\theta}) = \frac{1}{r} \sum_{r'=1}^r p^{(r')}(\boldsymbol{\theta})$ , and then simulating  $\mathbf{x} \sim p(\mathbf{x}|\boldsymbol{\theta})$ .}
  - 9:     For each pair in the mini-batch, sample  $0 < K - 1 < B$  contrasting parameters without replacement from the remainder of the mini-batch. {Equivalent to sampling each contrasting parameter independently from  $\tilde{p}(\boldsymbol{\theta})$ .}
  - 10:     Update  $\phi$  by stochastic gradient descent using objective  $\mathcal{L}(\phi) = -\frac{1}{B} \sum_{b=1}^B \log \frac{\exp(f_\phi(\boldsymbol{\theta}^{(b)}, \mathbf{x}^{(b)}))}{\sum_{k=1}^K \exp(f_\phi(\boldsymbol{\theta}^{(k)}, \mathbf{x}^{(b)}))}$ .
  - 11:   **end while**
  - 12:   Update posterior  $p^{(r)}(\boldsymbol{\theta}) \propto \exp(f_\phi(\boldsymbol{\theta}, \mathbf{x}_0))$ . {For SNPE-C, we can set  $p^{(r)}(\boldsymbol{\theta}) = q_\phi(\boldsymbol{\theta}|\mathbf{x}_0)$  directly.}
  - 13: **end for**
- 

where the progression from eq. (14) to eq. (15) is directly analogous to the progression from eq. (7) to eq. (8). Once more, we see that the optimal classifier implicitly learns the desired log-density ratio up to proportionality. Moreover, we still recover the correct ratio even if we sample parameters from a proposal  $\tilde{p}(\boldsymbol{\theta})$  instead of the prior  $p(\boldsymbol{\theta})$ , since  $\frac{\tilde{p}(\boldsymbol{\theta}|\mathbf{x})}{\tilde{p}(\boldsymbol{\theta})} \propto \frac{p(\boldsymbol{\theta}|\mathbf{x})}{p(\boldsymbol{\theta})}$ . With this result, the connection between SRE and SNPE-C becomes clear.

- Setting  $K = 2$  and parameterizing the classifier as a standard feed-forward neural network  $f_\phi(\boldsymbol{\theta}, \mathbf{x}) = \text{NN}_\phi(\boldsymbol{\theta}, \mathbf{x})$  recovers SRE. The derivation above also generalizes SRE from the binary case to multiple classes, and we will show this generally improves performance of the algorithm.
- The multi-class problem with parameterization of the classifier in terms of a density estimator  $f_\phi(\boldsymbol{\theta}, \mathbf{x}) = \log \frac{q_\phi(\boldsymbol{\theta}|\mathbf{x})}{p(\boldsymbol{\theta})}$  recovers SNPE-C. In this case,  $q_\phi(\boldsymbol{\theta}|\mathbf{x})$  yields a normalized approximation to the posterior when the density and the prior are supported on the same set (see appendix B for details), and may allow for efficient and exact evaluation and sampling depending on the particular choice of density estimator. This provides an alternative derivation to that given by Greenberg et al. (2019) of the correctness of the atomic variant of SNPE-C.

We believe this unifying perspective is beneficial to both practitioners and researchers, since comparison of these methods can now be cast in a more direct way, and the trade-offs of each more carefully considered. For instance, immediate questions arise concerning choice of contrasting set size, and whether reparameterizing the classifier in terms of a density estimator impacts inference. It is also interest-

ing to note how two lines of research into the likelihood-free problem converged essentially independently to very similar ideas; indeed, Hermans et al. (2020) compare to SNPE-C in their experiments without explicitly noting this connection. We present the general algorithm which encompasses both of these methods, iteratively refining a posterior approximation across a sequence of rounds, in algorithm 1.

**Connection to mutual information estimation.** Similar observations concerning density ratios have previously been made in the wider machine learning literature, but the implications for likelihood-free inference may not have been fully appreciated. The topic has received particular attention recently due to interest in mutual information estimation with neural networks (Belghazi et al., 2018; van den Oord et al., 2018; Poole et al., 2019). In this context, the classifier is often referred to as a critic, and multi-sample variational bounds have been derived for the mutual information, where it is also known that it can be advantageous to reparameterize the critic in terms of known densities. With this work, we seek to clarify these points within the likelihood-free inference community, and unify two promising lines of research which until now have been presented as distinct.

Of particular note is the connection of this contrastive learning scheme for likelihood-free inference to the multi-sample lower bound on the mutual information popularized by van den Oord et al. (2018). Imagine we sample a mini-batch of  $B$  pairs  $\{(\mathbf{x}^{(b)}, \boldsymbol{\theta}^{(b)})\}_{b=1}^B$  from our training set, where  $\boldsymbol{\theta}^{(b)} \sim \tilde{p}(\boldsymbol{\theta})$  was sampled from the some proposal, and  $\mathbf{x}^{(b)} \sim p(\mathbf{x}|\boldsymbol{\theta}^{(b)})$  was generated using the simulator. To generate a contrasting set of examples, we choose the remaining parameters in the batch, in effect sampling

marginally from the proposal  $\tilde{p}(\boldsymbol{\theta})$ . Indeed, this is the procedure used by SNPE-C to generate a contrasting set, and SRE can be extended similarly. Equivalently, this amounts to setting  $K = B$  in eq. (15). The classification loss for a minibatch is then

$$\mathcal{L}_B(\phi) = -\frac{1}{B} \sum_{b=1}^B \log \frac{\exp(f_\phi(\boldsymbol{\theta}^{(b)}, \mathbf{x}^{(b)}))}{\sum_{b'} \exp(f_\phi(\boldsymbol{\theta}^{(b')}, \mathbf{x}^{(b')})}. \quad (16)$$

The negative of this quantity is a well-known multi-sample lower bound on the mutual information between the two distributions defined by the joint  $\tilde{p}(\mathbf{x}, \boldsymbol{\theta}) = p(\mathbf{x} | \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\theta})$  and the product of marginals  $\tilde{p}(\mathbf{x}) \tilde{p}(\boldsymbol{\theta})$  (Poole et al., 2019), so that optimization of the classification loss can be equivalently interpreted as optimization of an estimate of the mutual information between these distributions.

This perspective gives some motivation for a sequential approach in likelihood-free methods, where in each round we sample parameters from a successively more accurate targeted proposal  $\tilde{p}(\boldsymbol{\theta})$  rather than the prior, which may be expected to decrease the mutual information across rounds. It also suggests that increasing the size of the contrasting set  $K$  may be beneficial, since this mutual information estimate is bounded above by  $\log K$ . However, neither SRE or SNPE-C needs to calculate the mutual information explicitly, and the density ratios can be learned for any  $K \geq 2$ .

## 4. Related Work

**Approximate Bayesian Computation (ABC).** We use the term ABC to describe the group of likelihood-free inference algorithms which require (i) choice of summary statistics  $\bar{\mathbf{x}}$  for observations  $\mathbf{x}$ , (ii) choice of a metric  $d(\cdot, \cdot)$  which computes distances between summary statistics, and (iii) choice of a tolerance  $\varepsilon > 0$  which determines whether two summary statistics are sufficiently similar based on the computed distance. Broadly, ABC methods draw approximate posterior samples by first proposing according to some scheme, and then rejecting samples whose summary statistics are distance more than  $\varepsilon$  from the observed summaries. Such algorithms have formed a central component of likelihood-free inference research for many years; see e.g. reviews by Beaumont (2010), Lintusaari et al. (2016) and Beaumont (2019). A drawback of ABC methods is that they generally require a large number of simulations to yield accurate results, which can be prohibitive when the simulator is expensive. Comparisons in the existing literature suggest that likelihood-free methods based on neural density estimators can outperform traditional approaches such as Sequential Monte Carlo ABC (Sisson et al., 2007; Beaumont et al., 2009; Toni et al., 2009) in terms of simulation cost, often requiring orders of magnitude fewer simulations (Papamakarios & Murray, 2016; Papamakarios et al., 2019b).

**Learning the likelihood.** Conversely to fitting a condi-

tional density estimator to the posterior directly, it is also possible to fit a synthetic likelihood to approximate  $p(\mathbf{x} | \boldsymbol{\theta})$ . A synthetic likelihood can be estimated separately for each parameter  $\boldsymbol{\theta}$  considered, e.g. as part of an MCMC scheme (Wood, 2010; Price et al., 2018; Fasiolo et al., 2018), or it can be amortized across parameters (Fan et al., 2013; Meeds & Welling, 2014). One advantage of the likelihood-learning approach is that parameters can be drawn from any proposal (Papamakarios et al., 2019b), and may even be chosen using active learning (Gutmann & Corander, 2016; Lueckmann et al., 2019; Järvenpää et al., 2018). In addition, the simulator may be such that learning the likelihood is easier than learning the posterior directly, which is often the case when the forward process is specified using simple probabilistic primitives, as in probabilistic programming (Gordon et al., 2014). Finally, a model of the likelihood is often easier to diagnose than a model of the posterior, since exact samples from  $p(\mathbf{x} | \boldsymbol{\theta})$  can be readily obtained (Dalmaso et al., 2019). On the other hand, any methods based on a surrogate likelihood require an additional inference step (such as MCMC) for posterior sampling, which adds complexity to the overall algorithm and can introduce further approximation error. Compounding this drawback is the fact that observations  $\mathbf{x}$  are often of much higher dimension than parameters  $\boldsymbol{\theta}$ , making conditional density estimation of observations a more difficult task. Indeed, Greenberg et al. (2019) show that the performance of the Sequential Neural Likelihood algorithm described by Papamakarios et al. (2019b) degrades when concatenating uninformative noise to observations  $\mathbf{x}$ .

**Contrastive learning.** The use of classification for density-ratio estimation is long-standing in machine learning, and is closely linked to the broad paradigm of learning-by-comparison, or contrastive learning. A prominent example of this line of thinking is Noise Contrastive Estimation (Gutmann & Hyvärinen, 2010), a method for fitting unnormalized distributions, or energy-based models, by learning relative probabilities through comparison with a known noise distribution. This approach forms the basis of the subsequent LFIRE method (Thomas et al., 2016). The branch of Generative Adversarial Networks literature concerned with fitting discriminators which approximate  $f$ -divergences (Goodfellow et al., 2014; Nowozin et al., 2016) has also made significant contribution to the ubiquity of density-ratio estimation techniques. Recently, contrastive learning has also shown great promise in unsupervised pre-training for vision tasks. A variant of Contrastive Predictive Coding (van den Oord et al., 2018; Hénaff et al., 2019) has demonstrated that a linear classifier on top of features learned in an unsupervised manner on ImageNet outperforms early convolutional networks in accuracy, while Momentum Contrast (He et al., 2019) outperforms supervised pre-training in a series of downstream detection and segmentation tasks.

**Density-ratio estimation for likelihood-free inference.**

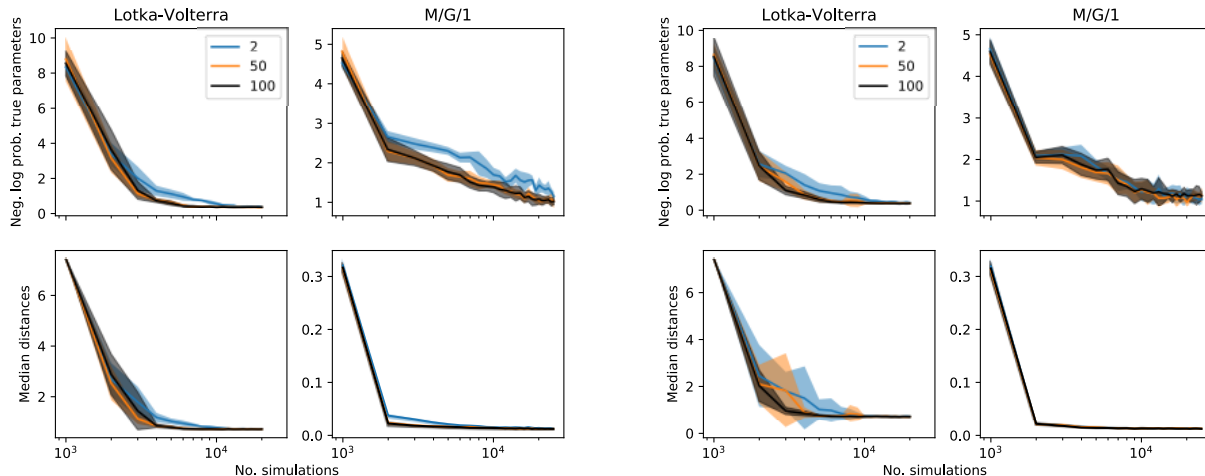


Figure 1: Comparison of SRE (left) and SNPE-C (right) across tasks and metrics, featuring mean and standard deviation across 10 random seeds. Negative log probability of true parameters is calculated using a kernel density estimate on posterior samples at the end of each round. Median distance uses the L2 distance between observations generated at each round and true observations. Increasing the contrasting-set size leads to more efficient inference, if only marginally in some cases.

The ratio of the likelihood  $p(\mathbf{x}|\theta)$  to the marginal  $p(\mathbf{x})$  is of central importance to Bayesian inference, and thus a natural quantity to estimate in likelihood-free inference. An early example of estimating this ratio is given by Izbicki et al. (2014); subsequent examples include LFIRE (Thomas et al., 2016) and SRE (Hermans et al., 2020), which we have already discussed. Another possibility is to learn the ratio between likelihoods  $p(\mathbf{x}|\theta_1)$  and  $p(\mathbf{x}|\theta_2)$ , with or without amortization across parameters  $\theta_1$  and/or  $\theta_2$ . This likelihood ratio can be used as the acceptance ratio in MCMC (Pham et al., 2014) or as a test statistic in hypothesis testing (Brehmer et al., 2018a), and can be learned by binary classification (Cranmer et al., 2016) or by regressing the joint-likelihood ratio if the latter is available (Brehmer et al., 2018b; Stoye et al., 2018). Other applications of density-ratio estimation in likelihood-free inference include: using binary-classification performance as a distance metric for ABC (Gutmann et al., 2018), estimating variational bounds (Tran et al., 2017), and testing goodness-of-fit of synthetic likelihoods (Dalmaso et al., 2019).

## 5. Experiments

In this section we compare SRE and SNPE-C experimentally, and discuss practical details of their implementation. For completeness we also consider Sequential Neural Likelihood (SNL, Papamakarios et al., 2019b), another sequential method for likelihood-free inference which fits a neural density estimator as a surrogate likelihood. The classifier for SRE is a feed-forward residual network with two residual blocks of 50 hidden units. All density estimators for SNL and SNPE-C use a MAF (Papamakarios et al., 2017) architecture, consisting of a stack of five MADEs (Germain

et al., 2015), each with two hidden layers of 50 units, and a standard normal base distribution. While MAF features density evaluation that is efficiently parallelizable across data dimensions, sampling is sequential, but the dimensionality of the parameter spaces considered is low enough that sampling is efficient. It is also worthwhile to point out the modularity of both SNL and SNPE-C, which readily allow for any off-the-shelf density estimator to be plugged in; the choice of MAF was made here in line with previous work.

We use axis-aligned slice sampling (Neal, 2003) as an MCMC method for SNL and SRE. A single MCMC chain persists across rounds for each method, where we perform burn-in of 200 iterations whenever the target distribution changes, and retain every tenth accepted sample. In each round, the parameters of each method are fit using stochastic gradient descent with the Adam (Kingma & Ba, 2015) optimizer, a learning rate of  $5e-4$ , and a mini-batch size of 100. To prevent overfitting, we perform early-stopping based on a held-out validation set of ten percent of the training data aggregated so far, stopping training when validation performance does not improve over 20 epochs. We compare likelihood-free methods using a testbed of three simulators, namely the Nonlinear Gaussian simulator with tractable likelihood described by Papamakarios et al. (2019b), along with the Lotka–Volterra predator-prey and M/G/1 queue models whose setups are detailed by Papamakarios & Murray (2016). For all tasks we acquire 1000 new simulations per round, running the Nonlinear Gaussian and M/G/1 tasks for 25 rounds, and Lotka–Volterra for 20. All experiments are repeated across 10 random seeds using a single GPU, and code is available at <https://github.com/conormdurkan/lfi>.

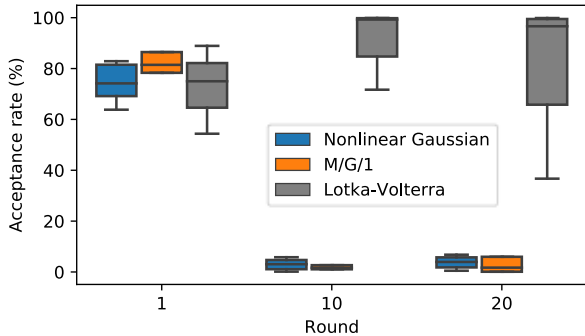


Figure 2: Distribution of rejection sampling acceptance rates for SNPE-C using  $10^7$  samples in rounds 1, 10, and 20 across 10 random seeds. Whiskers delineate 5<sup>th</sup> and 95<sup>th</sup> percentiles. Maximum likelihood training in the first round forces  $q_\phi$  to place density in the prior support. However, learning by classification in later rounds means there is no penalty for leaking significant density, and this phenomenon occurs in two of three tasks, with high variance in the third.

**A note on choice of prior.** A subtle detail arises when using SNPE-C with a density estimator whose support does not match that of the prior. Such scenarios arise commonly in practice, where simulators often impose hard constraints on parameter values, and a box uniform prior is chosen over an appropriate region. Indeed, each task considered in this section uses a box uniform prior on some region, in line with the setup in previous literature, while a MAF with a standard normal base distribution has support everywhere. However, this can lead to difficulties, as noted by Greenberg et al. (2019). Unlike maximum likelihood training, the classification objective defined by eq. (15) does not force  $q_\phi$  to place density in the prior support, but instead dictates only that density ratios are correct in this region; this is an inherent property of training a density estimator in an implicit fashion using a proxy classification task.

The consequence of this detail is that the density estimator can use an arbitrarily small portion of its density to fit the correct ratios if optimization using eq. (15) means it is advantageous to do so. This behaviour occurs commonly in practice, with posterior density ‘leaking’ from the prior support across rounds, such that acceptance rates for rejection sampling using the density estimator can become low. We stress, however, that this issue only arises with SNPE-C when the posterior support does not match that of the prior. On the other hand, there is no such similar notion for SRE, since the feed-forward network  $f_\phi$  defines a density only implicitly and is never evaluated outside of the prior support.

A full solution to this issue requires either reparameterizing the prior to be unconstrained, or alternatively reparameterizing the density estimator  $q_\phi$  to be supported only on the same region as the prior. We would argue that the former option is heavily influenced by domain-specific factors; it is

often natural for practitioners to prescribe hard constraints because simulators fail or run indefinitely for invalid values. For example, reparameterizing a strictly positive parameter on a log-scale may be problematic since exponentiated parameter values can lie arbitrarily close to zero, and although technically valid, may be numerically unstable. Empirically, we also found that box uniform priors tend to produce sharper posterior estimates than, say, Gaussian priors. On the other hand, reparameterizing the density estimator to be supported on some fixed range using an invertible squashing function, such as the sigmoid, leads to its own set of numerical issues. While theoretically valid, we encountered significant optimization difficulties when testing this approach, with densities often becoming degenerate due to numerical instabilities, and similar observations have been made in the density estimation literature (Durkan et al., 2019).

We mitigate this issue in practice by fitting SNPE-C in the first round using maximum likelihood, which is possible because we sample parameters from the prior initially. This forces the model to place density on the prior support in the first round. We also found the inductive bias of a MAF beneficial, since a series of elementwise affine transformations makes it difficult to transform a standard normal base distribution in an extreme way. In contrast, we found mixture models such as MDNs prone to entire mixture components drifting far from the prior support, while the remaining components are left to model the correct density ratios.

### 5.1. Size of the contrasting set

We first examine the role played by the size of the contrasting set. To this end, we run both SRE and SNPE-C for  $K \in \{2, 50, 100\}$ , and compare metrics across tasks. Results are shown in fig. 1. Across both multiple tasks and multiple metrics, the general trend suggests that increasing the size of the contrasting set results in more efficient inference, if only marginally in some cases. Computationally, setting  $K = 100$  means we are effectively evaluating mini-batches of size  $B^2$ , since the contrasting sets are generated using the full Cartesian product of  $B$  items in the original mini-batch. However, for  $B = 100$ , we found very little difference in wall-clock time for this operation due to efficient batch computation on the GPU. These experiments demonstrate the effectiveness of SRE in a relatively low-data regime; Hermans et al. (2020) use one million simulations for a majority of their tasks, whereas the performance of the same algorithm is tested here with orders of magnitude fewer. Based on these results, the rest of our experiments with SRE and SNPE-C are run using the full mini-batch with  $K = 100$  to generate the contrasting set.

### 5.2. Comparison of SRE and SNPE-C

The equivalence of SRE and SNPE-C also encourages investigation as to whether reparameterization of the classifier

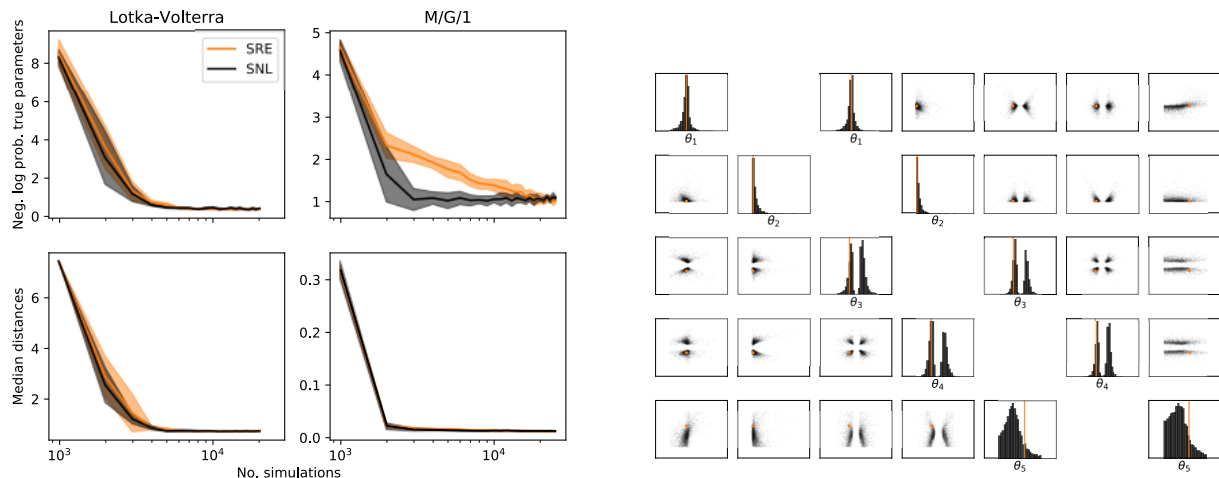


Figure 3: **Left:** Comparison of SRE and SNL across tasks and metrics, with negative log probability and median distance as in fig. 1. **Right:** Posterior samples for SNL (left) and SRE (right) on the Nonlinear Gaussian task; true parameters in orange.

in terms of a conditional density estimator has a noticeable impact on the quality of inference. Practitioners may be well-versed in the use of standard classifiers from the modern machine learning toolbox, but less familiar with the recent surge of flexible density estimators based on neural networks. Comparison of these approaches may thus guide practical use, since if no discrepancy in performance is observed, SNPE-C has a clear advantage in that a posterior approximation with density evaluation and i.i.d. sampling is available at the end of training. We again refer to fig. 1 for results, which indeed show SNPE-C closely matching SRE. As a control, we also ran SNPE-C with MCMC sampling, and observed essentially the same performance as with i.i.d. sampling. However, it is important to consider these results in the context of fig. 2, which visualize rejection sampling acceptance rates across each task for a number of rounds. As discussed, SNPE-C leaks density outside of the prior support, with acceptance rates around 1% meaning i.i.d. sampling becomes up to 100 times slower. Nevertheless, this may still be appreciably faster than MCMC sampling; using a neural density estimator, we may generate large batches consisting of thousands of i.i.d. samples in parallel with a few passes of a neural network, while MCMC requires at least one neural network pass to evaluate a transition kernel which may not even result in a single effective sample.

### 5.3. Comparison of SRE and SNL

Finally we compare SRE with SNL. Like SRE, SNL is also reliant on MCMC sampling to generate posterior samples in each round. Results are displayed in fig. 3. Overall, SNL seems to outperform SRE, particularly on the M/G/1 task. Though we don’t do so here, SNL also makes it possible to carry out goodness-of-fit testing to check correctness of its surrogate likelihood model, since sampling observations from the model is straightforward, which is not the case for

SRE. However, as mentioned already, if observations are high-dimensional, SRE has a clear advantage in that it need not parameterize a distribution over this high-dimensional space, but only take as input this object to the classifier.

## 6. Discussion

Recent work has demonstrated cases where likelihood-free inference methods that leverage deep representations can reduce the number of simulations needed for accurate and efficient inference by orders of magnitude. However, there are several open problems: in their current implementations, the methods described in this paper still (i) require too many simulations if a simulator takes many days to run, (ii) have difficulty exploiting multiple observations, and (iii) don’t deal gracefully with failed simulations. Moreover, the methods in the literature have used a variety of density estimation and density-ratio estimation principles and architectures. As a result, practitioners facing a new problem domain will have difficulty choosing from the alternatives.

In this work, we identify a strong connection between two strands of the literature, exemplified by SRE and SNPE-C, that were previously compared as unrelated methods. Our connection contributes new understanding of these existing approaches, and shows that SRE can and should be generalized to use multi-class classification.

It would be easy to assume that SNPE-C is harder to fit because it uses a conditional density estimator rather than a classifier as in SRE. However, our control experiments show that — when trained using the same codebase and classification objective — the representations give similar results. In such cases, using a conditional density estimator should be preferred, because it gives a normalized posterior and removes the need to use MCMC. However, we also iden-



tify open practical problems with SNPE-C for priors with compact support, which mean that MCMC may be necessary in the case of extremely high rejection rates. We hope that our comparison helps clarify the underlying choices for likelihood-free practitioners, and will motivate researchers to progress the open problems that we have identified.

## Acknowledgements

The authors would like to thank Theophane Weber for helpful feedback. This work was supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh.

## References

- Alsing, J., Charnock, T., Feeney, S., and Wandelt, B. Fast likelihood-free cosmology with neural density estimators and active learning. *Monthly Notices of the Royal Astronomical Society*, 488(3):4440–4458, 2019.
- Beaumont, M. A. Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41(1):379–406, 2010.
- Beaumont, M. A. Approximate Bayesian computation. *Annual Review of Statistics and Its Application*, 6(1):379–403, 2019.
- Beaumont, M. A., Cornuet, J.-M., Marin, J.-M., and Robert, C. P. Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990, 2009.
- Belghazi, M. I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. Mutual information neural estimation. *International Conference on Machine Learning*, 2018.
- Bishop, C. M. Mixture density networks. Technical Report NCRG/94/004, Aston University, 1994.
- Brehmer, J., Cranmer, K., Louppe, G., and Pavez, J. Constraining effective field theories with machine learning. *Physical Review Letters*, 121(11):111801, 2018a.
- Brehmer, J., Louppe, G., Pavez, J., and Cranmer, K. Mining gold from implicit models to improve likelihood-free inference. *arXiv preprint arXiv:1805.12244*, 2018b.
- Cranmer, K., Pavez, J., and Louppe, G. Approximating likelihood ratios with calibrated discriminative classifiers. *arXiv preprint arXiv:1506.02169*, 2016.
- Cranmer, K., Brehmer, J., and Louppe, G. The frontier of simulation-based inference. *arXiv preprint arXiv:1911.01429*, 2019.
- Dalmaso, N., Lee, A. B., Izbicki, R., Pospisil, T., and Lin, C.-A. Validation of approximate likelihood and emulator models for computationally intensive simulations. *arXiv preprint arXiv:1905.11505*, 2019.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Cubic-spline flows. *Workshop on Invertible Neural Networks and Normalizing flows at the International Conference on Machine Learning*, 2019.
- Fan, Y., Nott, D. J., and Sisson, S. A. Approximate Bayesian computation via regression density estimation. *Stat*, 2(1): 34–48, 2013.
- Fasiolo, M., Wood, S. N., Hartig, F., and Bravington, M. V. An extended empirical saddlepoint approximation for intractable likelihoods. *Electronic Journal of Statistics*, 12(1):1544–1578, 2018.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. MADE: Masked autoencoder for distribution estimation. *International Conference on Machine Learning*, 2015.
- Gonçalves, P. J., Lueckmann, J.-M., Deistler, M., Nonnenmacher, M., Öcal, K., Bassetto, G., Chintaluri, C., Podlaski, W. F., Haddad, S. A., Vogels, T. P., Greenberg, D. S., and Macke, J. H. Training deep neural density estimators to identify mechanistic models of neural dynamics. *bioRxiv*, pp. 838383, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 2014.
- Gordon, A. D., Henzinger, T. A., Nori, A. V., and Rajamani, S. K. Probabilistic programming. *Proceedings of the on Future of Software Engineering*, pp. 167–181, 2014.
- Greenberg, D. S., Nonnenmacher, M., and Macke, J. H. Automatic posterior transformation for likelihood-free inference. *International Conference on Machine Learning*, 2019.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. *International Conference on Artificial Intelligence and Statistics*, 2010.
- Gutmann, M. U. and Corander, J. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(125): 1–47, 2016.
- Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. Likelihood-free inference via classification. *Statistics and Computing*, 28(2):411–425, 2018.

- Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York, 2001.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- Hénaff, O. J., Razavi, A., Doersch, C., Eslami, S., and van den Oord, A. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019.
- Hermans, J., Begy, V., and Louppe, G. Likelihood-free MCMC with amortized approximate ratio estimators. *International Conference on Machine Learning*, 2020.
- Izbicki, R., Lee, A. B., and Schafer, C. High-dimensional density ratio estimation with extensions to approximate likelihood computation. *International Conference on Artificial Intelligence and Statistics*, 2014.
- Järvenpää, M., Gutmann, M., Vehtari, A., and Marttinen, P. Efficient acquisition rules for model-based approximate Bayesian computation. *Bayesian Analysis*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Lintusaari, J., Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, 66(1):e66–e82, 2016.
- Lueckmann, J.-M., Goncalves, P. J., Bassetto, G., Öcal, K., Nonnenmacher, M., and Macke, J. H. Flexible statistical inference for mechanistic models of neural dynamics. *Advances in Neural Information Processing Systems*, 2017.
- Lueckmann, J.-M., Bassetto, G., Karaletsos, T., and Macke, J. H. Likelihood-free inference with emulator networks. *Symposium on Advances in Approximate Bayesian Inference*, 2019.
- Meeds, E. and Welling, M. GPS-ABC: Gaussian process surrogate approximate Bayesian computation. *Conference on Uncertainty in Artificial Intelligence*, 2014.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Neal, R. M. Slice sampling. *Annals of Statistics*, pp. 705–741, 2003.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. *Advances in Neural Information Processing Systems*, 2016.
- Papamakarios, G. *Neural density estimation and likelihood-free inference*. PhD thesis, School of Informatics, University of Edinburgh, 2019.
- Papamakarios, G. and Murray, I. Fast  $\varepsilon$ -free inference of simulation models with Bayesian conditional density estimation. *Advances in Neural Information Processing Systems*, 2016.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems*, 2017.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019a.
- Papamakarios, G., Sterratt, D. C., and Murray, I. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. *International Conference on Artificial Intelligence and Statistics*, 2019b.
- Pham, K. C., Nott, D. J., and Chaudhuri, S. A note on approximating ABC-MCMC using flexible classifiers. *Stat*, 3(1):218–227, 2014.
- Poole, B., Ozair, S., van den Oord, A., Alemi, A. A., and Tucker, G. On variational bounds of mutual information. *International Conference on Machine Learning*, 2019.
- Price, L. F., Drovandi, C. C., Lee, A., and Nott, D. J. Bayesian synthetic likelihood. *Journal of Computational and Graphical Statistics*, 27(1):1–11, 2018.
- Sisson, S. A., Fan, Y., and Tanaka, M. M. Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007.
- Sisson, S. A., Fan, Y., and Beaumont, M. *Handbook of approximate Bayesian computation*. Chapman and Hall/CRC, 2018.
- Stoye, M., Brehmer, J., Louppe, G., Pavez, J., and Cranmer, K. Likelihood-free inference with an improved cross-entropy estimator. *arXiv preprint arXiv:1808.00973*, 2018.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density ratio estimation in machine learning*. Cambridge University Press, 2012. ISBN 0521190177.
- Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. Likelihood-free inference by ratio estimation. *arXiv preprint arXiv:1611.10242*, 2016.
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A., and Stumpf, M. P. H. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical

systems. *Journal of The Royal Society Interface*, 6(31): 187–202, 2009.

Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. *Advances in Neural Information Processing Systems*, 2017.

Uribe, B., Côté, M.-A., Gregor, K., Murray, I., and Larochelle, H. Neural autoregressive distribution estimation. *Journal of Machine Learning Research*, 17(205): 1–37, 2016.

van den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

Wood, S. N. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.