

---

# Adversarial Variational Optimization of Non-Differentiable Simulators

---

Gilles Louppe  
University of Liège

Joeri Hermans  
University of Liège

Kyle Cranmer  
New York University

## Abstract

Complex computer simulators are increasingly used across fields of science as generative models tying parameters of an underlying theory to experimental observations. Inference in this setup is often difficult, as simulators rarely admit a tractable density or likelihood function. We introduce Adversarial Variational Optimization (AVO), a likelihood-free inference algorithm for fitting a non-differentiable generative model incorporating ideas from generative adversarial networks, variational optimization and empirical Bayes. We adapt the training procedure of generative adversarial networks by replacing the differentiable generative network with a domain-specific simulator. We solve the resulting non-differentiable minimax problem by minimizing variational upper bounds of the two adversarial objectives. Effectively, the procedure results in learning a proposal distribution over simulator parameters, such that the JS divergence between the marginal distribution of the synthetic data and the empirical distribution of observed data is minimized. We evaluate and compare the method with simulators producing both discrete and continuous data.

## 1 Introduction

In fields of science such as particle physics, epidemiology or population genetics, computer simulators are used to describe complex data generation processes. These simulators relate observations  $\mathbf{x}$  to the parameters  $\theta$  of an underlying theory or mechanistic model.

Simulators are usually specified as procedural implementations of forward, stochastic processes involving latent variables  $\mathbf{z}$ . Rarely do these simulators admit a tractable density, or likelihood,  $p(\mathbf{x}|\theta)$ . The prevalence and significance of this problem has motivated an active research effort in so-called *likelihood-free inference* algorithms such as Approximate Bayesian Computation (ABC) and density estimation-by-comparison algorithms (Beaumont et al., 2002; Marjoram et al., 2003; Sisson et al., 2007; Sisson and Fan, 2011; Marin et al., 2012; Cranmer et al., 2015).

In parallel, with the introduction of variational autoencoders (Kingma and Welling, 2013) and generative adversarial networks (Goodfellow et al., 2014), there has been a vibrant research program around implicit generative models based on neural networks (Mohamed and Lakshminarayanan, 2016). While some of these models also do not admit a tractable density, they are all differentiable by construction. In addition, generative models based on neural networks are highly parameterized and the model parameters have no obvious interpretation. In contrast, scientific simulators can be thought of as highly regularized generative models as they typically have few parameters and they are endowed with some level of interpretation. In this setting, inference on the model parameters  $\theta$  is often of more interest than the latent variables  $\mathbf{z}$ .

In this work, we introduce Adversarial Variational Optimization (AVO), a likelihood-free inference algorithm for non-differentiable, implicit generative models. We adapt the adversarial training procedure of generative adversarial networks by replacing the implicit generative network with a domain-based scientific simulator, and solve the resulting non-differentiable minimax problem by minimizing variational upper bounds of the adversarial objectives. The objective of the algorithm is to match the marginal distribution of the synthetic data to the empirical distribution of observations. Therefore, AVO enables empirical Bayes in the likelihood-free setting – which should not be mistaken with (Bayesian) posterior inference.

## 2 Problem statement

We consider a family of parameterized densities  $p(\mathbf{x}|\boldsymbol{\theta})$  defined implicitly through the simulation of a stochastic generative process, where  $\mathbf{x} \in \mathbb{R}^d$  is the data and  $\boldsymbol{\theta}$  are the parameters of interest. The simulation may involve some complicated latent process where  $\mathbf{z} \in \mathcal{Z}$  is a latent variable providing an external source of randomness. Unlike implicit generative models defined by neural networks, we do not assume  $\mathbf{z}$  to be a fixed-size vector with a simple density. Instead, the dimension of  $\mathbf{z}$  and the nature of its components (uniform, normal, discrete, continuous, etc.) are inherited from the control flow of the simulation code and may depend on  $\boldsymbol{\theta}$  in some intricate way. Moreover, the dimension of  $\mathbf{z}$  may be much larger than the dimension of  $\mathbf{x}$ .

We assume that the stochastic generative process that defines  $p(\mathbf{x}|\boldsymbol{\theta})$  is specified through a non-differentiable deterministic function  $g(\cdot; \boldsymbol{\theta}) : \mathcal{Z} \rightarrow \mathbb{R}^d$ . Operationally,

$$\mathbf{x} \sim p(\mathbf{x}|\boldsymbol{\theta}) \triangleq \mathbf{z} \sim p(\mathbf{z}|\boldsymbol{\theta}), \mathbf{x} = g(\mathbf{z}; \boldsymbol{\theta}) \quad (1)$$

such that the density  $p(\mathbf{x}|\boldsymbol{\theta})$  can be written as

$$p(\mathbf{x}|\boldsymbol{\theta}) = \int_{\{\mathbf{z}: g(\mathbf{z}; \boldsymbol{\theta}) = \mathbf{x}\}} p(\mathbf{z}|\boldsymbol{\theta}) \mu(d\mathbf{z}), \quad (2)$$

where  $\mu$  is a probability measure.

Given some observed data  $\{\mathbf{x}_i | i = 1, \dots, N\}$  drawn from the (unknown) true distribution  $p_r(\mathbf{x})$ , our goal is to estimate the parameters  $\boldsymbol{\theta}^*$  that minimize some divergence or some distance  $\rho$  between  $p_r(\mathbf{x})$  and the implicit model  $p(\mathbf{x}|\boldsymbol{\theta})$ . That is,

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \rho(p_r(\mathbf{x}), p(\mathbf{x}|\boldsymbol{\theta})). \quad (3)$$

## 3 Background

### 3.1 Generative adversarial networks

Generative adversarial networks (GANs) were first proposed by Goodfellow et al. (2014) as a way to build an implicit generative model capable of producing samples from random noise  $\mathbf{z}$ . The core principle of GANs is to pit a generative model  $g(\cdot; \boldsymbol{\theta})$  against an adversarial classifier  $d(\cdot; \boldsymbol{\phi}) : \mathbb{R}^d \rightarrow [0, 1]$  that has for antagonistic objective to recognize real data  $\mathbf{x}$  from generated data  $\tilde{\mathbf{x}} = g(\mathbf{z}; \boldsymbol{\theta})$ . Both models  $g$  and  $d$  are trained simultaneously, in such a way that  $g$  learns to fool its adversary  $d$  (which happens when  $g$  produces samples comparable to the observed data), while  $d$  continuously adapts to changes in  $g$ .

In practice, the discriminator  $d$  and the generator  $g$  are usually trained with alternating stochastic gradient

descent in order to respectively minimize

$$\begin{aligned} \mathcal{L}_d(\boldsymbol{\phi}) &= \mathbb{E}_{\mathbf{x} \sim p_r(\mathbf{x})} [-\log(d(\mathbf{x}; \boldsymbol{\phi}))] \\ &\quad + \mathbb{E}_{\tilde{\mathbf{x}} \sim p(\mathbf{x}|\boldsymbol{\theta})} [-\log(1 - d(\tilde{\mathbf{x}}; \boldsymbol{\phi}))] \quad (4) \\ \mathcal{L}_g(\boldsymbol{\theta}) &= \mathbb{E}_{\tilde{\mathbf{x}} \sim p(\mathbf{x}|\boldsymbol{\theta})} [\log(1 - d(\tilde{\mathbf{x}}; \boldsymbol{\phi}))], \quad (5) \end{aligned}$$

where  $\mathcal{L}_d$  corresponds to the binary cross-entropy between true and synthetic data and  $\mathcal{L}_g$  is the negative of  $\mathcal{L}_d$  restricted to synthetic data.

When  $d$  is trained to optimality before each (infinitesimally small) parameter update of the generator, it can be shown that the original adversarial learning procedure of Goodfellow et al. (2014) amounts to minimizing the Jensen-Shannon divergence JSD between the distributions  $p_r(\mathbf{x})$  and  $p(\mathbf{x}|\boldsymbol{\theta})$ . Of course this assumption is never met in practice and it is often observed that the GAN alternating optimization procedure does not lead to convergence. As a result, recent research has focused on finding better training algorithms (e.g., Salimans et al., 2016; Metz et al., 2016; Arjovsky et al., 2017; Gulrajani et al., 2017; Roth et al., 2017) for GANs, as well as gaining a better theoretical understanding of the training dynamics (e.g., Arjovsky and Bottou, 2017; Mescheder et al., 2017a; Nagarajan and Kolter, 2017). In this work, we follow Mescheder et al. (2018) and adapt the GAN training procedure by adding a regularization term

$$R_1(\boldsymbol{\phi}) = \mathbb{E}_{\mathbf{x} \sim p_r(\mathbf{x})} [|\|\nabla_{\boldsymbol{\phi}} d(\mathbf{x}; \boldsymbol{\phi})\|^2] \quad (6)$$

to the loss  $\mathcal{L}_d$  of the discriminator. Under suitable assumptions, this regularization term guarantees the (local) convergence of the training procedure, while keeping the original GAN algorithm otherwise unchanged.

### 3.2 Variational optimization

Variational optimization (Staines and Barber, 2012; Staines and Barber, 2013) and evolution strategies (Wierstra et al., 2011) are general optimization techniques that can be used to form a differentiable bound on the optima of a non-differentiable function. Given a function  $f$  to minimize, these techniques are based on the observation that

$$\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) \leq \mathbb{E}_{\boldsymbol{\theta} \sim q(\boldsymbol{\theta}|\boldsymbol{\psi})} [f(\boldsymbol{\theta})] = U(\boldsymbol{\psi}), \quad (7)$$

where  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$  is a proposal distribution with parameters  $\boldsymbol{\psi}$  over input values  $\boldsymbol{\theta}$ . That is, the minimum of a set of function values is always less than or equal to any of their average. Provided that the proposal distribution is flexible enough, the parameters  $\boldsymbol{\psi}$  can be updated to place its mass arbitrarily tight around the optimum  $\boldsymbol{\theta}^* = \min_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta})$ .

Under mild restrictions outlined by Staines and Barber (2012), the bound  $U(\boldsymbol{\psi})$  is differentiable with respect

to  $\psi$ , and using the log-likelihood trick its gradient can be rewritten as:

$$\begin{aligned}
 \nabla_{\psi} U(\psi) &= \nabla_{\psi} \mathbb{E}_{\theta \sim q(\theta|\psi)} [f(\theta)] \\
 &= \nabla_{\psi} \int q(\theta|\psi) f(\theta) d\theta \\
 &= \int \nabla_{\psi} q(\theta|\psi) f(\theta) d\theta \\
 &= \int q(\theta|\psi) \nabla_{\psi} \log q(\theta|\psi) f(\theta) d\theta \\
 &= \mathbb{E}_{\theta \sim q(\theta|\psi)} [\nabla_{\psi} \log q(\theta|\psi) f(\theta)] \quad (8)
 \end{aligned}$$

Effectively, this means that provided that the score function  $\nabla_{\psi} \log q(\theta|\psi)$  of the proposal is known and that one can evaluate  $f(\theta)$  for any  $\theta$ , then one can construct empirical estimates of Eqn. 8, which can in turn be used to minimize  $U(\psi)$  with stochastic gradient descent (or a variant thereof, robust to noise and parameter scaling).

In reinforcement learning, Eqn. 8 similarly appears in the context of policy gradients, where  $f(\theta)$  corresponds to a reward signal for the action  $\theta$  and  $q(\theta|\psi)$  corresponds to a policy  $\pi_{\psi}$  that we aim to optimize. In this context, empirical estimates of Eqn. 8 are better known as REINFORCE estimates (Williams, 1992).

## 4 Adversarial variational optimization

### 4.1 Algorithm

The alternating stochastic gradient descent on  $\mathcal{L}_d$  and  $\mathcal{L}_g$  in GANs (Section 3.1) implicitly assumes that the generator  $g$  is a differentiable function. In the setting where we are interested in estimating the parameters of a fixed non-differentiable simulator (Section 2) – as opposed to learning the generative model itself – gradients  $\nabla_{\theta} g$  either do not exist or are not accessible. As a result, gradients  $\nabla_{\theta} \mathcal{L}_g$  cannot be constructed and the optimization procedure cannot be carried out.

In this work, we propose to rely on variational optimization to minimize  $\mathcal{L}_d$  and  $\mathcal{L}_g$ , thereby bypassing the non-differentiability of  $g$ . We consider a proposal distribution  $q(\theta|\psi)$  over the parameters of the simulator  $g$  and alternately minimize the variational upper bounds

$$U_d(\phi) = \mathbb{E}_{\theta \sim q(\theta|\psi)} [\mathcal{L}_d(\phi)] \quad (9)$$

$$U_g(\psi) = \mathbb{E}_{\theta \sim q(\theta|\psi)} [\mathcal{L}_g(\theta)] \quad (10)$$

respectively over  $\phi$  and  $\psi$ . The discriminator  $d$  is therefore no longer pit against a single generator  $g$ , but instead against a hierarchical family of generators induced by the proposal distribution.

When updating the discriminator parameters  $\phi$ , unbiased estimates of  $\nabla_{\phi} U_d$  can be obtained by directly

evaluating the gradient of  $U_d$  over mini-batches of real and synthetic data. When updating the proposal parameters  $\psi$ ,  $\nabla_{\psi} U_g$  can be estimated as described in the previous section with  $f(\theta) = \mathcal{L}_g(\theta)$ . That is,

$$\nabla_{\psi} U_g = \mathbb{E}_{\substack{\theta \sim q(\theta|\psi), \\ \tilde{\mathbf{x}} \sim p(\mathbf{x}|\theta)}} [\nabla_{\psi} \log q(\theta|\psi) \log(1 - d(\tilde{\mathbf{x}}; \phi))], \quad (11)$$

which we can approximate with mini-batches of synthetic data.

While the latter REINFORCE-like gradient estimator is unbiased, it is well known that it also suffers from high variance, which makes the optimization unstable and difficult. A common remedy to this issue (Williams, 1992) is to make use of the fact that

$$\begin{aligned}
 &\mathbb{E}_{\theta \sim q(\theta|\psi)} [\nabla_{\psi} \log q(\theta|\psi) f(\theta)] \\
 &= \mathbb{E}_{\theta \sim q(\theta|\psi)} [\nabla_{\psi} \log q(\theta|\psi) (f(\theta) - b)] \quad (12)
 \end{aligned}$$

for any constant  $b$ . The choice of the baseline  $b$  does not bias the gradient estimator, but it can however have an effect on its variance. For AVO, we pick the baseline which minimizes the variance of the empirical estimates of  $\nabla_{\psi} U_g$ , that is

$$b = \frac{\mathbb{E} [(\nabla_{\psi} \log q(\theta|\psi))^2 (1 - \log(d(\tilde{\mathbf{x}}; \phi))^2)]}{\mathbb{E} [(\nabla_{\psi} \log q(\theta|\psi))^2]}. \quad (13)$$

For completeness, Algorithm 1 outlines the proposed Adversarial Variational Optimization (AVO) procedure, as built on top of GAN with  $R_1$  regularization.

### 4.2 Parameter Point Estimates

The variational objectives 9-10 effectively replace the modeled data distribution of Eqn. 1 with the parameterized marginal distribution of the generated data

$$q(\mathbf{x}|\psi) = \int p(\mathbf{x}|\theta) q(\theta|\psi) d\theta. \quad (14)$$

We can think of  $q(\mathbf{x}|\psi)$  as a *variational program* as described by Ranganath et al. (2016), though more complicated than a simple reparameterization of normally distributed noise  $\mathbf{z}$  through a differentiable function. In our case, the variational program is a marginalized, non-differentiable simulator. Its density is intractable; nevertheless, it can generate samples for  $\mathbf{x}$  whose expectations are differentiable with respect to  $\psi$ . Operationally, we sample from this marginal model via

$$\mathbf{x} \sim q(\mathbf{x}|\psi) \triangleq \theta \sim q(\theta|\psi), \mathbf{z} \sim p(\mathbf{z}|\theta), \mathbf{x} = g(\mathbf{z}; \theta). \quad (15)$$

We can view the optimization of  $q(\mathbf{x}|\psi)$  with respect to  $\psi$  through the lens of empirical Bayes, where the data are used to optimize a prior within the family  $q(\theta|\psi)$ .

---

**Algorithm 1** Adversarial variational optimization (AVO).
 

---

**Inputs:** Observed data  $\{\mathbf{x}_i \sim p_r(\mathbf{x})\}_{i=1}^N$ , simulator  $g$ .  
**Outputs:** Proposal distribution  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$ , such that  $q(\mathbf{x}|\boldsymbol{\psi}) \approx p_r(\mathbf{x})$ .  
**Hyper-parameters:** The number  $k$  of training iterations of the discriminator  $d$  (default:  $k = 1$ ),  
 The size  $M$  of a mini-batch (default:  $M = 32$ ),  
 The  $R_1$  regularization coefficient  $\lambda$  (default:  $\lambda = 10$  (Mescheder et al., 2018)),  
 The entropy penalty coefficient  $\gamma$  (default:  $\gamma = 0$ ).  
 The baseline strategy  $b$  in REINFORCE estimates (default: Eqn. 13).

- 1:  $q(\boldsymbol{\theta}|\boldsymbol{\psi}) \leftarrow$  prior on  $\boldsymbol{\theta}$  (with differentiable and known density)
- 2: **while**  $\boldsymbol{\psi}$  has not converged **do**
- 3:     **for**  $i = 1$  to  $k$  **do** ▷ Update  $d$
- 4:         Sample true data  $\{\mathbf{x}_m \sim p_r(\mathbf{x}), y_m = 1\}_{m=1}^{M/2}$ .
- 5:         Sample synthetic data  $\{\boldsymbol{\theta}_m \sim q(\boldsymbol{\theta}|\boldsymbol{\psi}), \mathbf{z}_m \sim p(\mathbf{z}|\boldsymbol{\theta}_m), \tilde{\mathbf{x}}_m = g(\mathbf{z}_m; \boldsymbol{\theta}_m), y_m = 0\}_{m=M/2+1}^M$ .
- 6:          $\nabla_{\phi} U_d \leftarrow \frac{1}{M} \sum_{m=1}^M \nabla_{\phi} [-y_m \log(d(\mathbf{x}_m; \phi)) - (1 - y_m) \log(1 - d(\mathbf{x}_m; \phi))]$
- 7:          $\nabla_{\phi} R_1 \leftarrow \frac{1}{M/2} \sum_{m=1}^{M/2} \nabla_{\phi} [|\nabla_{\phi} d(\mathbf{x}_m; \phi)|^2]$
- 8:          $\phi \leftarrow \text{RMSPROP}(\nabla_{\phi} U_d + \lambda \nabla_{\phi} R_1)$
- 9:     **end for**
- 10:     Sample synthetic data  $\{\boldsymbol{\theta}_m \sim q(\boldsymbol{\theta}|\boldsymbol{\psi}), \mathbf{z}_m \sim p(\mathbf{z}|\boldsymbol{\theta}_m), \tilde{\mathbf{x}}_m = g(\mathbf{z}_m; \boldsymbol{\theta}_m)\}_{m=1}^M$ . ▷ Update  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$
- 11:      $\nabla_{\psi} U_g \leftarrow \frac{1}{M} \sum_{m=1}^M [\nabla_{\psi} \log q(\boldsymbol{\theta}_m|\boldsymbol{\psi})(\log(1 - d(\tilde{\mathbf{x}}_m; \phi)) - b)]$
- 12:      $\nabla_{\psi} H \leftarrow \frac{1}{M} \sum_{m=1}^M \nabla_{\psi} [-q(\boldsymbol{\theta}_m|\boldsymbol{\psi}) \log q(\boldsymbol{\theta}_m|\boldsymbol{\psi})]$
- 13:      $\boldsymbol{\psi} \leftarrow \text{RMSPROP}(\nabla_{\psi} U_g + \gamma \nabla_{\psi} H)$
- 14: **end while**

---

Since the GAN procedure effectively minimizes the Jensen-Shannon divergence between  $p_r(\mathbf{x})$  and  $q(\mathbf{x}|\boldsymbol{\psi})$ ,  $\boldsymbol{\psi}^*$  corresponds with the maximum marginal likelihood estimator advocated by Rubin (1984). When the model is well specified,  $\boldsymbol{\psi}^*$  coincides with the true data-generating parameter; however, if the model is misspecified, the  $\boldsymbol{\psi}^*$  is typically different from the maximum likelihood estimator (MLE). Thus, if the simulator  $p(\mathbf{x}|\boldsymbol{\theta})$  is misspecified,  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$  will attempt to smear it so that the marginal model  $q(\mathbf{x}|\boldsymbol{\psi})$  is closer to  $p_r(\mathbf{x})$ . However, if the simulator is well specified, then  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$  will concentrate its mass around the true data-generating parameter.

In order to more effectively target point estimates  $\boldsymbol{\theta}^*$ , we can also augment Eqn. 10 with an entropic regularization term  $H(q(\boldsymbol{\theta}|\boldsymbol{\psi}))$ , such that

$$U_g = \mathbb{E}_{\boldsymbol{\theta} \sim q(\boldsymbol{\theta}|\boldsymbol{\psi})}[\mathcal{L}_g] + \gamma H(q(\boldsymbol{\theta}|\boldsymbol{\psi})), \quad (16)$$

where  $\gamma \in \mathbb{R}^+$  is a hyper-parameter controlling the trade-off between the generator objective and the tightness of the proposal distribution and  $H$  is the Shannon entropy. For small values of  $\gamma$ , proposal distributions with large entropy are not penalized, which results in learning a smeared variation of the original simulator. On the other hand, for large values of  $\gamma$ , the procedure is encouraged to fit a proposal distribution with low entropy, which has the effect of concentrating its density tightly around one or a few  $\boldsymbol{\theta}$  values.

Finally, we note that very large penalties may eventually make the optimization unstable, as the variance of  $\nabla_{\psi} \log q(\boldsymbol{\theta}_m|\boldsymbol{\psi})$  typically increases as the entropy of the proposal decreases. Depending on the proposal, it may also be possible to always arbitrarily decrease the entropy, without necessarily producing accurate parameter estimates. In this case, properly controlling for  $\gamma$  and the number of training epochs is critical. For this reason, in all but the first experiment hereafter, entropic regularization is disabled by default ( $\gamma = 0$ ).

## 5 Experiments

### 5.1 Illustrative example

As a first illustrative experiment, we evaluate inference for a discrete Poisson distribution with unknown mean  $\lambda$ . We artificially consider the distribution as a parameterized simulator, from which we can only generate data.

The discrete observed data is sampled from a Poisson with mean  $\lambda^* = 7$ . Algorithm 1 is run for 3000 iterations with mini-batches of size  $M = 32$  and the following configuration. For the discriminator  $d$ , we use a 3-layer MLP with 20 hidden nodes per layer and PReLU activation units. For estimating  $\lambda^*$ , we parameterize  $\boldsymbol{\theta}$  as  $\log(\lambda)$  and use a univariate Gaussian proposal distribution  $q(\boldsymbol{\theta}|\boldsymbol{\psi})$  initialized with a mean of

$\log(1)$  and a variance of  $0.5^2$ . The  $R_1$  regularization coefficient is set to 10, and the entropy penalty is evaluated at both  $\gamma = 0$  and  $\gamma = 0.0001$ . The learning rate of RMSPROP is set to 0.001, both for the discriminator  $d$  and the proposal  $q$ .

The top left plot in Figure 1 illustrates the resulting proposal distributions  $q(\theta|\psi)$  after running AVO. For both  $\gamma = 0$  and  $\gamma = 0.0001$ , the proposal distributions correctly concentrate their density around the true parameter value  $\log(\lambda^*) = 1.94$ . Under the effect of entropic regularization, the proposal distribution for  $\gamma = 0.0001$  concentrates its mass more tightly, yielding in this case more precise inference. The top right plot compares the model distributions to the true distribution. As theoretically expected from adversarial training, we see that the resulting distributions align with the true distribution, with in this case visually slightly better results for the penalized model. The bottom plot of Figure 1 evaluates the negative log-likelihood of the true parameters  $\lambda^*$  with respect to the number of simulated samples. For the two settings, the loss steadily decreases as the proposal converges towards the nominal parameter value. This short example highlights that adversarial variational optimization works despite the discreteness of the data and the lack of access to the density  $p(\mathbf{x}|\theta)$  or its gradient.

## 5.2 High-energy particle collisions

As a more challenging example, we now turn to a particle physics inference problem. We consider the PYTHIA simulator (Sjöstrand et al., 2008) for high-energy particle collisions routinely used by physicists at the Large Hadron Collider. We simulate electron-positron collisions at a center-of-mass energy of 91.2 GeV, in which a  $Z$  boson is produced and decays to quarks. We assume a detector that emulates a  $32 \times 32$  spherical uniform grid in pseudorapidity  $\eta$  and in azimuthal angle  $\phi$ , covering  $(\eta, \phi) \in [-5, 5] \times [0, 2\pi]$ . The detector is parameterized by an offset parameter  $\theta$  in the  $z$ -axis relative to the beam crossing point (Borisyyak, 2018). An offset of  $\theta = 0$  means that the sphere is centered at the collision point, while  $\theta = 1$  leads to a shift of roughly one pixel.

The inference problem we are interested in is the estimation of the offset parameter  $\theta$  from a set of  $32 \times 32$ -dimensional observations. This task is representative of calibration and alignment tasks, which are critical in experimental particle physics as they have significant impact on the accuracy of reconstruction algorithms.

The leftmost plots of Figure 2 show the average detector response for two distinct offsets  $\theta = 0$  and  $\theta = 1$ . The remaining plots illustrate individual random samples for these configurations. The figures highlight the




Figure 1: Discrete Poisson model with unknown mean. (Top left) Proposal distributions  $q(\theta|\psi)$  after training. For both  $\gamma = 0$  and  $\gamma = 0.0001$ , the distributions correctly concentrate their density around the true value  $\log(\lambda^*)$ . Entropic regularization ( $\gamma = 0.0001$ ) results in a tighter density. (Top right) Model distributions  $q(\mathbf{x}|\psi)$  after training. This plot shows that the resulting parameterizations of the simulator closely reproduce the true distribution. (Bottom) Negative log-likelihood  $-q(\theta^*|\psi)$  of the target parameters, as function of the number of simulated samples.




Figure 2: PYTHIA-alignment: Samples. (Top row): Average detector response (approximated over 200,000 samples) for the detector offset  $\theta = 0$ , along with 3 individual random samples. (Bottom row) Same but for the detector offset  $\theta = 1$ . Given the sparsity and variability of the simulated events, these plots highlight the difficulty in observing a difference between samples from one or the other parameter setting.

challenging difficulty of the inference problem: the difference between the average responses is barely noticeable, while individual samples are very sparse and reflect a wide range of variability. These samples also stress the critical role of a relevant summary statistic on such high-dimensional data, which is required in likelihood-free inference methods such as ABC.

For this experiment, we consider observed data simulated at the nominal value  $\theta^* = 1$ . Algorithm 1 is run for 5000 iterations with all hyper-parameters set to their default values. The discriminator  $d$  is defined as a 9-layer MLP with 600 hidden nodes per layer and PReLU activations. The proposal distribution is initialized as a Gaussian with zero mean and unit variance. As shown in the top plot of Figure 3, the proposal distribution properly centers around the target  $\theta^* = 1$  after training. The bottom plots in the figure also illustrate the convergence of AVO as a function of the number of simulations. Despite the complexity of the PYTHIA simulator, the sparsity, variability and high-dimensionality of the observations, as well as the absence of any domain knowledge, AVO properly identifies the target parameter within a reasonable number of simulations. As suggested clearly by the bottom right plot of Figure 3, where the negative log-likelihood  $-\log q(\theta^*|\psi)$  has not yet converged, more accurate results could certainly be obtained by running AVO for more iterations. Equivalently, we anticipate room for hyper-parameter tuning.

Finally, let us also comment on the bump around  $10^5$  simulations in the left bottom plot of Figure 3. This illustrates the particular scenario in which a temporary deviation in the mean of the proposal from the target parameter value is compensated by the variance of the proposal, which thereby results in even lower negative log-likelihood  $-\log q(\theta^*|\psi)$ . This is confirmed by the right bottom plot, where no such bump is observed.

### 5.3 Benchmarks

**Methods.** In this section, we systematically evaluate AVO on benchmark problems. We compare our algorithm against ABC-SMC (Toni and Stumpf, 2009) and BOLFI (Gutmann and Corander, 2016) as baselines. ABC-SMC is the most commonly used instance of Approximate Bayesian Computation. It makes use of importance sampling to improve efficiency. BOLFI is a simulation-efficient likelihood-free inference algorithm that combines Bayesian optimization with a Gaussian process-based approximation of the likelihood of summary statistics of the data. For ABC-SMC and BOLFI, we respectively use the PYABC (Klinger et al., 2017) and the ELFI (Lintusaari et al., 2018) implementations. All hyper-parameters are set to the default values recommended in these packages.

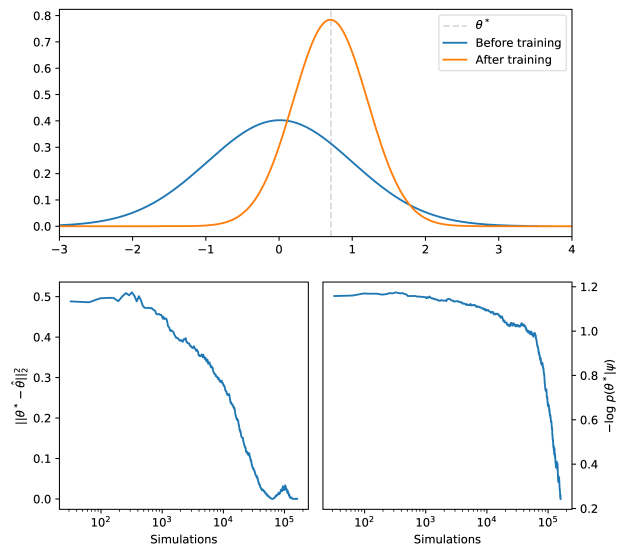


Figure 3: PYTHIA-alignment: training. (Top) Proposal distribution  $q(\theta|\psi)$  before and after AVO. (Bottom left). Distance between the target parameter  $\theta^*$  and the average parameter under the proposal, as a function of the number of simulated samples. (Bottom right). Negative log-likelihood of the target parameter value  $\theta^*$  under the proposal, as a function of the number of simulated samples.

**Inference tasks.** We evaluate all three methods on four inference tasks of increasing difficulty. These tasks include discrete, continuous, low-dimensional and high-dimensional observations. For each task, we evaluate the quality of inference in terms of squared error for 15 different target parameter values  $\theta_i^*$ , for  $i = 1, \dots, 15$ . For each target value, we consider a data set with 100,000 observations representing  $\mathbf{x} \sim p_r(\mathbf{x})$ . All methods evaluated share the same simulation budget (160,000 samples).

- *Poisson.* This inference problem is the same as in Section 5.1, with  $\theta_i^* \sim \mathcal{U}(0, 4)$ . The discriminator  $d$  is defined as a 3-layer MLP with PReLU activation units and 600 nodes per hidden layer.
- *CARL-Multidimensional.* We reproduce the inference problem originally defined in Section 4.2 of (Cranmer et al., 2015). The generator is parameterized by two parameters  $\alpha$  and  $\beta$  and produces 5-dimensional continuous data  $\mathbf{x} \in \mathbb{R}^5$ . For our benchmark, we consider  $\alpha_i^* \sim \mathcal{U}(-2, 2)$  and  $\beta_i^* \sim \mathcal{U}(-2, 2)$ . The discriminator  $d$  is defined as a 4-layer MLP with PReLU activation units and 100 nodes per hidden layer.
- *Weinberg.* We consider a simplified simulator for electron-positron collisions, as described in Appendix A.1. We consider  $E_i^{\text{beam}} \sim \mathcal{U}(43, 47)$  and

$G_i^f \sim \mathcal{U}(0, 2)$ . The discriminator  $d$  is defined as 4-layer MLP with PReLU activation units and 1000 nodes per hidden layer.

- *PYTHIA-alignment*. This inference problem is the same as in Section 5.2, with  $\theta_i^* \sim \mathcal{U}(-1.5, 1.5)$ . The discriminator  $d$  is defined as a 9-layer MLP with PReLU activation units and 600 nodes per hidden layer.

The summary statistics used in ABC-SMC and BOLFI are the same. For Poisson, CARL-Multidimensional and Weinberg, the summary statistics are defined as the Euclidean distance between (the bins of) the histogram of the observations generated at  $\theta_i^*$  and (the bins of) the histogram of simulated data. For PYTHIA-alignment, the summary statistics is defined as the  $\ell_2$  norm between the average image of the observed data and the average image of the simulated samples at the model parameter. In both methods, 128 simulation samples are generated per model parameter evaluation. The priors used are identical to the uniform priors used for generating the 15 problems  $\theta_i^*$ . For AVO, the proposal distribution is initialized as a Gaussian of zero mean and unit variance.

In contrast to some related works, we focus on the setting where we have more than one observation  $\mathbf{x}$  from the data distribution  $p_r(\mathbf{x})$ . For this reason, we do not consider likelihood-free benchmarks such as the M/G/1 queue model, the Lotka-Volterra population model or the Hodgkin-Huxley neuron model, which are all defined as inference problems from single observations. We anticipate that AVO is less appropriate for this use case, as the discriminator  $d$  would not be expected to provide a good learning signal for fitting the simulator parameters. The proper treatment of this scenario is left as future work.

**Results.** Figure 4 summarizes our results for AVO, ABC-SMC and BOLFI on each of the four inference problems. Each plot reports as a box plot the squared distance of the best fit  $\hat{\theta}_i$  to the nominal model parameters  $\theta_i^*$ , for  $i = 1, \dots, 15$ . In AVO,  $\hat{\theta}_i$  corresponds to the mode of the final proposal. For ABC-SMC and BOLFI,  $\hat{\theta}_i$  is the maximum a posteriori estimate under the posterior. Best fit values are comparable since we assume uniform priors.

The figure clearly indicates AVO works better on average compared to ABC-SMC and BOLFI. We attribute this superior performance primarily to the fact that AVO is not limited by the deficiencies of a hand-crafted summary statistic. Instead, AVO benefits from a high-capacity discriminator that dynamically adapts to the inference problem and to the current proposal. This is clearly apparent for PYTHIA-alignment, where

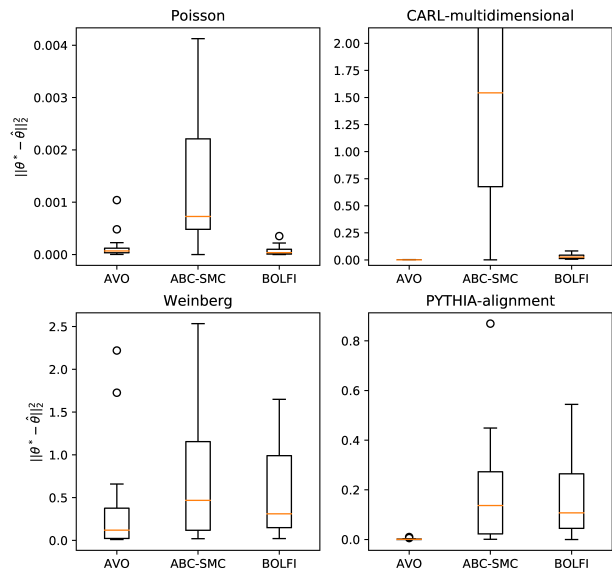


Figure 4: Benchmark results comparing AVO against ABC-SMC and BOLFI. AVO shows superior performance against these methods. This advantage stems from the fact that AVO is not limited by the suboptimality of an ad hoc summary statistic. Instead, the discriminator in AVO can dynamically adapt to the inference problem.

a generic summary statistic leads to a sub-optimal estimator. By contrast, because of the high-capacity discriminator  $d$ , AVO has no issue in guiding the proposal towards a solution, despite the high-dimensionality of the observations or the complexity of underlying generative process. Of course, ABC-SMC and BOLFI can be improved by engineering better summary statistics, but this requires a deep understanding of the problem. While it is not illustrated here, the active learning strategy of BOLFI shows better sample efficiency than AVO, in the sense that it can often reach a good solution within a smaller simulation budget. Finally, for the Weinberg benchmark, we observe that there is no clear winner in terms of the squared error. This is mainly due to the an approximate degeneracy between the parameters leading to a very broad minimum and a number of solutions that fit the observed data distribution (see Appendix A.2).

## 6 Related work

This work sits at the intersection of several lines of research related to likelihood-free inference, approximate Bayesian computation (ABC), implicit generative models, and variational inference. Viewed from the literature around implicit generative models based on neural networks, the proposed method

can be considered as a direct adaptation of generative adversarial networks (Goodfellow et al., 2014) to non-differentiable simulators using variational optimization (Staines and Barber, 2012). From the point of view of likelihood-free inference, where non-differentiable simulators are the norm, our contributions are threefold. First is the process of lifting the expectation with respect to the non-differentiable simulator  $\mathbb{E}_{\tilde{\mathbf{x}} \sim p(\mathbf{x}|\boldsymbol{\theta})}$  to a differentiable expectation with respect to the variational program  $\mathbb{E}_{\tilde{\mathbf{x}} \sim q(\mathbf{x}|\boldsymbol{\psi})}$ . Secondly, is the introduction of a novel form of variational inference that works in a likelihood-free setting. Thirdly, AVO can be viewed as a form of empirical Bayes where the prior is optimized based on the data.

As for many likelihood-free inference algorithms, AVO is intimately tied to a class of algorithms that can be framed as density estimation-by-comparison, as reviewed in (Mohamed and Lakshminarayanan, 2016). In most cases, these inference algorithms are formulated as an iterative two-step process where the model distribution is first compared to the true data distribution and then updated to make it more comparable to the latter. Relevant work in this direction includes those that rely on a classifier to estimate the discrepancy between the observed data and the model distributions (Gutmann and Hyvärinen, 2012; Cranmer et al., 2015, 2016; Dutta et al., 2016; Gutmann et al., 2017; Rosca et al., 2017). Of direct relevance in the likelihood-free setup, Hamiltonian ABC (Meeds et al., 2015) estimates gradients with respect to  $\boldsymbol{\theta}$  through finite differences from multiple forward passes of the simulator with variance reduction strategies based on controlling the source of randomness used for the latent variable  $\mathbf{z}$ . Sharing similar foundational principles as AVO but developed independently, the SPIRAL algorithm (Ganin et al., 2018) makes use of the Wasserstein GAN objective and variants of REINFORCE gradient estimates to adversarially train an agent that synthesizes programs controlling a non-differentiable graphics engine in order to reconstruct target images or perform unconditional generation.

Likewise, AVO closely relates to recent extensions of GANs, such as ALI (Dumoulin et al., 2016), BiGANs (Donahue et al., 2016),  $\alpha$ -GAN (Rosca et al., 2017), AVB (Mescheder et al., 2017b), and the PC-Adv algorithm of (Huszár, 2017), which add an inference network to the generative model. Each of these assume a tractable density  $p(\mathbf{x}|\boldsymbol{\theta})$  that is differentiable with respect to  $\boldsymbol{\theta}$ , which is not satisfied in the likelihood-free setting. Our lifting of the non-differentiable simulator  $p(\mathbf{x}|\boldsymbol{\theta})$  to the variational program  $q(\mathbf{x}|\boldsymbol{\psi})$  provides the ability to differentiate expectations with respect to  $\boldsymbol{\psi}$  as in Eqn 8; however, the density  $q(\mathbf{x}|\boldsymbol{\psi})$  is still intractable. Moreover, we do not attempt to define

a recognition model  $q(\mathbf{z}, \boldsymbol{\theta}|\boldsymbol{\psi})$  as the latent space  $\mathcal{Z}$  of many real-world simulators is complicated and not amenable to a neural recognition model.

This work has also many connections to work on variational inference, in which the goal is to optimize the recognition model  $q(\mathbf{z}, \boldsymbol{\theta}|\boldsymbol{\psi})$  so that it is close to the true posterior  $p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{x})$ . There have been efforts to extend variational inference to intractable likelihoods; however, many require restrictive assumptions. In (Tran et al., 2017), the authors consider Variational Bayes with an Intractable Likelihood (VBIL). In that approach “the only requirement is that the intractable likelihood can be estimated unbiasedly.” In the case of simulators, they propose to use the ABC-likelihood with an  $\epsilon$ -kernel. The ABC likelihood is only unbiased as  $\epsilon \rightarrow 0$ , thus this method inherits the drawbacks of the ABC-likelihood including the choice of summary statistics and the inefficiency in evaluating the ABC likelihood for high-dimensional data and small  $\epsilon$ . More recently, (Tran et al., 2017) adapted variational inference to hierarchical implicit models defined on simulators. In this work, the authors step around the intractable likelihoods by reformulating the optimization of the ELBO in terms of a neural and differentiable approximation  $r$  of the log-likelihood ratio  $\log \frac{p}{q}$ , thereby effectively using the same core principle as used in GANs (Mohamed and Lakshminarayanan, 2016). With a similar objective, (McCarthy et al., 2017) adapt variational inference to a non-differentiable cardiac simulator by maximizing the ELBO using Bayesian optimization, hence bypassing altogether the need for gradient estimates.

## 7 Summary

In this work, we develop a likelihood-free inference algorithm for non-differentiable, implicit generative models. The algorithm combines adversarial training with variational optimization to minimize variational upper bounds on the otherwise non-differentiable adversarial objectives. The AVO algorithm enables empirical Bayes through variational inference in the likelihood-free setting. This approach does not incur the inefficiencies of an ABC-like rejection sampler nor the disadvantages of likelihood-free inference algorithms that rely on ad hoc summary statistics. When the model is well-specified, the AVO algorithm provides point estimates for the generative model, which asymptotically corresponds to the data generating parameters. Experimental results highlight the good performance of AVO in comparison to the well-established ABC-SMC and BOLFI algorithms.



## Acknowledgments

We would like to thank Lukas Heinrich for helpful comments regarding the electron–positron annihilation simulation. We would also like to thank Rajesh Ranganath and Dustin Tran for enlightening discussions and feedback. GL and KL were both supported through NSF ACI-1450310 at the time of the research, additionally KC is supported through PHY-1505463 and PHY-1205376. JH acknowledges the financial support from F.R.S-FNRS for his FRIA PhD scholarship.

## References

- Agostinelli, S. et al. (2003). GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303.
- Alwall, J., Herquet, M., Maltoni, F., Mattelaer, O., and Stelzer, T. (2011). MadGraph 5 : Going Beyond. *JHEP*, 06:128.
- Arjovsky, M. and Bottou, L. (2017). Towards Principled Methods for Training Generative Adversarial Networks. *ArXiv e-prints*.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein GAN. *ArXiv e-prints*.
- Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035.
- Borisov, M. (2018). Pythia-mill. <https://github.com/maxim-borisov/pythia-mill>.
- Cranmer, K., Pavez, J., and Louppe, G. (2015). Approximating likelihood ratios with calibrated discriminative classifiers. *arXiv preprint arXiv:1506.02169*.
- Cranmer, K., Pavez, J., Louppe, G., and Brooks, W. (2016). Experiments using machine learning to approximate likelihood ratios for mixture models. In *Journal of Physics: Conference Series*, volume 762, page 012034. IOP Publishing.
- Donahue, J., Krähenbühl, P., and Darrell, T. (2016). Adversarial feature learning. *arXiv preprint arXiv:1605.09782*.
- Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., and Courville, A. (2016). Adversarially learned inference. *arXiv preprint arXiv:1606.00704*.
- Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. (2016). Likelihood-free inference by ratio estimation. *ArXiv e-prints*.
- Ganin, Y., Kulkarni, T., Babuschkin, I., Eslami, S. M. A., and Vinyals, O. (2018). Synthesizing Programs for Images using Reinforced Adversarial Learning. *ArXiv e-prints*.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. (2017). Improved Training of Wasserstein GANs. *ArXiv e-prints*.
- Gutmann, M. U. and Corander, J. (2016). Bayesian optimization for likelihood-free inference of simulator-based statistical models. *The Journal of Machine Learning Research*, 17(1):4256–4302.
- Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. (2017). Likelihood-free inference via classification. *Statistics and Computing*, pages 1–15.
- Gutmann, M. U. and Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13(Feb):307–361.
- Huszár, F. (2017). Variational Inference using Implicit Distributions. *ArXiv e-prints*.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *CoRR*, abs/1312.6114.
- Klinger, E., Rickert, D., and Hasenauer, J. (2017). pyabc: distributed, likelihood-free inference. *bioRxiv*.
- Lintusaari, J., Vuollekoski, H., Kangasrsi, A., Skytn, K., Järvenp, M., Marttinen, P., Gutmann, M., Vehtari, A., Corander, J., and Kaski, S. (2018). Elfi: Engine for likelihood free inference.
- Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate bayesian computational methods. *Statistics and Computing*, pages 1–14.
- Marjoram, P., Molitor, J., Plagnol, V., and Tavaré, S. (2003). Markov chain monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26):15324–15328.
- McCarthy, A., Rodriguez, B., and Minchole, A. (2017). Variational Inference over Non-differentiable Cardiac Simulators using Bayesian Optimization. *ArXiv e-prints*.
- Meeds, E., Leenders, R., and Welling, M. (2015). Hamiltonian abc. *arXiv preprint arXiv:1503.01916*.
- Mescheder, L., Geiger, A., and Nowozin, S. (2018). Which training methods for gans do actually converge? In *International Conference on Machine Learning*, pages 3478–3487.
- Mescheder, L., Nowozin, S., and Geiger, A. (2017a). The numerics of gans. In *Advances in Neural Information Processing Systems*, pages 1825–1835.

- Mescheder, L. M., Nowozin, S., and Geiger, A. (2017b). Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. *CoRR*, abs/1701.04722.
- Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. (2016). Unrolled Generative Adversarial Networks. *ArXiv e-prints*.
- Mohamed, S. and Lakshminarayanan, B. (2016). Learning in Implicit Generative Models. *ArXiv e-prints*.
- Nagarajan, V. and Kolter, J. Z. (2017). Gradient descent gan optimization is locally stable. In *Advances in Neural Information Processing Systems*, pages 5585–5595.
- Ranganath, R., Altosaar, J., Tran, D., and Blei, D. M. (2016). Operator Variational Inference. *ArXiv e-prints*.
- Rosca, M., Lakshminarayanan, B., Warde-Farley, D., and Mohamed, S. (2017). Variational approaches for auto-encoding generative adversarial networks. *arXiv preprint arXiv:1706.04987*.
- Roth, K., Lucchi, A., Nowozin, S., and Hofmann, T. (2017). Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems*, pages 2018–2028.
- Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Ann. Statist.*, 12(4):1151–1172.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242.
- Sisson, S. A. and Fan, Y. (2011). *Likelihood-free MCMC*. Chapman & Hall/CRC, New York.[839].
- Sisson, S. A., Fan, Y., and Tanaka, M. M. (2007). Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765.
- Sjöstrand, T., Mrenna, S., and Skands, P. (2008). A brief introduction to pythia 8.1. *Computer Physics Communications*, 178(11):852–867.
- Staines, J. and Barber, D. (2012). Variational Optimization. *ArXiv e-prints*.
- Staines, J. and Barber, D. (2013). Optimization by variational bounding. In *ESANN 2013 proceedings, 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pages 473–478.
- Toni, T. and Stumpf, M. P. (2009). Simulation-based model selection for dynamical systems in systems and population biology. *Bioinformatics*, 26(1):104–110.
- Tran, D., Ranganath, R., and Blei, D. M. (2017). Hierarchical Implicit Models and Likelihood-Free Variational Inference. *ArXiv e-prints*.
- Tran, M.-N., Nott, D. J., and Kohn, R. (2017). Variational bayes with intractable likelihood. *Journal of Computational and Graphical Statistics*, (just-accepted).
- Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., and Schmidhuber, J. (2011). Natural Evolution Strategies. *ArXiv e-prints*.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.