

---

# Neural Posterior Estimation with Differentiable Simulators

---

Justine Zeghal<sup>\*1</sup> François Lanusse<sup>\*2</sup> Alexandre Boucaud<sup>\*1</sup> Benjamin Remy<sup>2</sup> Eric Aubourg<sup>3</sup>

## Abstract

Simulation-Based Inference (SBI) is a promising Bayesian inference framework that alleviates the need for analytic likelihoods to estimate posterior distributions. Recent advances using neural density estimators in SBI algorithms have demonstrated the ability to achieve high-fidelity posteriors, at the expense of a large number of simulations ; which makes their application potentially very time-consuming when using complex physical simulations. In this work we focus on boosting the sample-efficiency of posterior density estimation using the gradients of the simulator. We present a new method to perform Neural Posterior Estimation (NPE) with a differentiable simulator. We demonstrate how gradient information helps constrain the shape of the posterior and improves sample-efficiency.

## 1. Introduction

Various scientific fields use complex computer simulations to describe real physical processes as accurately as possible. These simulators map the parameter space  $\theta$  to a simulation  $x$  through an implicit likelihood  $p(x|\theta)$  that makes the inverse process of constraining the parameter space from data intractable in practice. Simulation-Based Inference (SBI), also known as Likelihood-free inference, provides a framework to alleviate that problem in different ways.

One can approximate the likelihood distribution  $p(x|\theta)$  (Wood, 2010; Papamakarios et al., 2018; Lueckmann et al., 2018), or the likelihood ratio  $r(x, \theta', \theta_t) = \frac{p(x|\theta')}{p(x|\theta_t)}$  from the simulations (Cranmer et al., 2015; Izbicki et al., 2014; Thomas et al., 2016; Durkan et al., 2020), and use a sampling method to estimate the posterior. Others choose an

amortized method by directly approximating the posterior distribution  $p(\theta|x)$  (Blum & François, 2009; Papamakarios & Murray, 2018; Lueckmann et al., 2017; Greenberg et al., 2019). But such methods treat the simulator as a black-box implicit distribution by considering only forward simulations and discarding all information on the internal process.

Brehmer et al. (2020) proposed to work with augmented data such as the gradients of the simulator and introduced a way to approximate the likelihood distribution and the likelihood ratio which leverages this augmented data and thus improves sample efficiency and inference quality.

In this work, we extend the work of Brehmer et al. (2020) and propose the first Neural Posterior Estimation method augmented with gradients of the simulator. Implementing this approach necessitates the use of a particular kind of Normalizing Flows, called Smooth Normalizing Flows (Köhler et al., 2021), which have the property of having well defined, smooth, and expressive gradients. We apply our approach on a standard SBI benchmark problem (Lotka-Volterra) and recover acceleration factors when using simulation gradients consistent with the results obtained with the NLE and NRE methods of Brehmer et al. (2020).

## 2. Simulation-Based Inference with Differentiable Simulators

Bayesian inference aims to infer the parameters  $\theta_0$  that have generated a given observation  $x_0$ . From Bayes theorem we have

$$p(\theta|x_0) = \frac{p(x_0|\theta)p(\theta)}{p(x_0)} \propto p(x_0|\theta)p(\theta), \quad (1)$$

and since  $x_0$  is the result of a large number of transformations involving a large number of latent variables  $z$ , the marginal likelihood

$$p(x_0|\theta) = \int p(x_0|\theta, z)p(z)dz, \quad (2)$$

is intractable.

SBI is particularly useful in this case since it provides a framework to approximate the posterior without using an analytic likelihood. To approximate  $p(\theta|x_0)$  SBI algorithms

---

<sup>\*</sup>Equal contribution <sup>1</sup>Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France <sup>2</sup>Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191, Gif-sur-Yvette, France <sup>3</sup>Université Paris Cité, CNRS, CEA, Astroparticule et Cosmologie, F-75013 Paris, France. Correspondence to: Justine Zeghal <zeghal@apc.in2p3.fr>.

require the observation  $x_0$ , a prior  $p(\theta)$  for the model parameters  $\theta$  and a simulator  $x \sim p(x|\theta)$  to sample from the intractable likelihood.

In this work we are interested in doing NPE which aims to directly learn the following distribution

$$p(\theta|x) \propto \int p(x|\theta, z)p(z|\theta)p(\theta)dz, \quad (3)$$

using a NDE such as a Normalizing Flow (hereafter NF). One can train the NDE to learn the approximated distribution  $p_\varphi(\theta|x)$  from samples  $(x, \theta)$  of the joint distribution by minimizing

$$\mathcal{L}_{\text{NLL}} = \mathbb{E}_{p(\theta, x)} [-\log p_\varphi(\theta|x)], \quad (4)$$

and then evaluate it on the given observation  $x_0$  to get the approximated posterior  $p_\varphi(\theta|x = x_0) \approx p(\theta|x_0)$ .

As recognized in Brehmer et al. (2020), when the simulators are differentiable one can extract for each simulation the gradients w.r.t. simulation parameters, which provides significantly more information than the samples from the simulator, and can be used to help constrain the posterior density estimates obtained by SBI. Therefore, in our case, we extract for each simulation  $(x_i, \theta_i)$  the gradient of the joint log-probability of the simulator with respect to input parameters. In the following, we will note this gradient as

$$\begin{aligned} \nabla_\theta \log p(\theta|x, z) &= \nabla_\theta \log p(x|\theta, z) \\ &\quad + \nabla_\theta \log p(z|\theta) \\ &\quad + \nabla_\theta \log p(\theta), \end{aligned} \quad (5)$$

where  $z$  are latent stochastic variables of the simulator. In a slight abuse of language, we will refer to  $\nabla_\theta \log p(\theta|x, z)$  as the *joint score*, but it should be noted that the conventional definition of the score function in statistics (also adopted in Brehmer et al. (2020)) is the gradient of the log likelihood  $\nabla_\theta \log p(x|\theta, z)$ .

We can then define a direct score matching loss

$$\mathcal{L}_{\text{SM}} = \mathbb{E}_{p(x, z, \theta)} [\|\nabla_\theta \log p(\theta|x, z) - \nabla_\theta \log p_\varphi(\theta|x)\|_2^2]. \quad (6)$$

Inspired by Brehmer et al. (2020), this quantity is minimized by

$$\begin{aligned} &\mathbb{E}_{p(z|x, \theta)} [\nabla_\theta \log p(\theta|x, z)] \\ &= \mathbb{E}_{p(z|x, \theta)} \left[ \nabla_\theta \log \frac{p(\theta, z|x)p(x)}{p(z, x)} \right] \\ &= \mathbb{E}_{p(z|x, \theta)} [\nabla_\theta \log p(\theta, z|x)] \\ &= \nabla_\theta \log p(\theta|x) \end{aligned} \quad (7)$$

meaning that by minimizing  $\mathcal{L}_{\text{SM}}$  we approximate the intractable marginal score.

Ultimately, we train our NDE using the following combined loss:

$$\mathcal{L} = \mathcal{L}_{\text{NLL}} + \lambda \mathcal{L}_{\text{SM}}, \quad (8)$$

where  $\lambda$  is a hyper-parameter which may be used to tune the score contribution to the loss function. The optimal value for  $\lambda$  typically depends on the problem considered.

Note that in order to train the NDE by constraining its score  $\nabla_\theta \log p_\varphi(\theta|x)$ , it needs to be sufficiently smooth with respect to  $\theta$ , motivating the development of dedicated NF architectures described in the next section.

While in this work we focus on NPE, which allows for amortized inference, as demonstrated in Brehmer et al. (2020) similar equations can be formulated for enhancing likelihood and likelihood ratio estimation with gradient information.

### 3. Smooth Normalizing Flows

To learn the approximated distribution  $p_\varphi(\theta|x)$  we focus on NFs (Rezende & Mohamed, 2015), a class of density estimators parameterized by neural networks that provides tractable density estimation. The key idea of NFs is to transform a simple density distribution (e.g. a multivariate Normal distribution) through a series of bijective functions to reconstruct a complex target distribution. Several strategies exist to define these bijections, here we focus on models relying on coupling transforms (Dinh et al., 2014; 2016)  $\phi: \mathbb{R}^D \mapsto \mathbb{R}^D$  of the form:

$$\begin{cases} x_{1:d} &= z_{1:d} \\ x_{d+1:D} &= g_{z_{1:d}}^\varphi(z_{d+1:D}) \end{cases} \quad (9)$$

where  $g_{z_{1:d}}^\varphi$  is an invertible *coupling function*, which is tuned based on the first  $d$  dimensions of the input vector  $z$ . Standard models like the RealNVP (Dinh et al., 2016) use affine coupling functions  $g_y^\varphi(z) = \sigma_\varphi(y) \cdot z + \mu_\varphi(y)$  with  $\sigma_\varphi$  and  $\mu_\varphi$  are neural networks, and  $\sigma_\varphi$  strictly positive.

To be trained on both simulations and gradients using the combined loss (8), a NF requires expressive derivatives  $\nabla_\theta \log p_\varphi(\theta|x)$  (6). If the coupling function is not smooth enough, i.e. does not have well defined and non-trivial higher-order gradients, the NF cannot be trained by score matching. We illustrate this in the case of the affine coupling of a RealNVP on the middle panel of Figure 1, which fails to train under a score matching loss by lack of expressivity.

To overcome this issue, Köhler et al. (2021) propose a coupling function based on a  $C^\infty$  diffeomorphism on  $]0, 1[$

$$f_y^\varphi(x) = \frac{\sigma(x) - \sigma(x_0)}{\sigma(x_1) - \sigma(x_0)} \cdot (1 - c_\varphi(y)) + c_\varphi(y) \cdot x \quad (10)$$

$$\sigma(x) = c_\varphi(y) \cdot x + \frac{1 - c_\varphi(y)}{1 + \exp(-\rho(x))} \quad (11)$$

$$\rho(x) = a_\varphi(y) \cdot \left( \log \left( \frac{x}{1-x} \right) + b_\varphi(y) \right) \quad (12)$$

with  $[x_0, x_1] \subset ]0, 1[$ , and  $a_\varphi, b_\varphi, c_\varphi$  learned using a smooth Neural Network (NN).

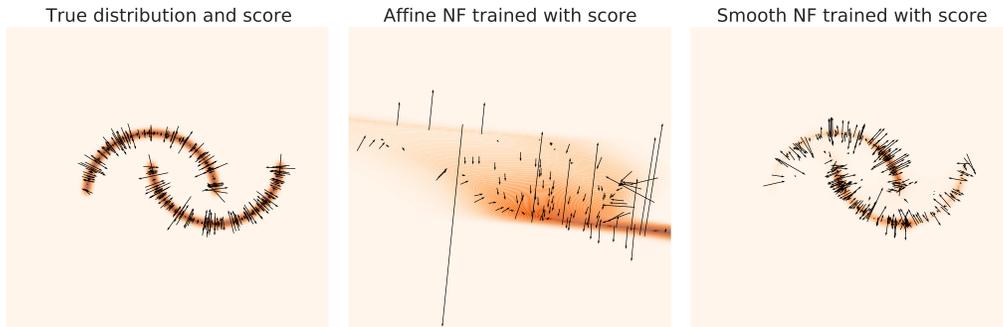


Figure 1. Two-moons experiment: learning a bi-modal distribution using score matching loss only. Left: Two-moons true distribution (orange) and score (arrows). Middle: distribution and score approximated with a standard affine RealNVP (Dinh et al., 2016). Right: distribution and score approximated with a smooth NF (Köhler et al., 2021). It illustrates that Affine coupling layers cannot be trained using a score matching loss.

A single sigmoidal function like the one proposed here is not very expressive, but the sum of  $C^\infty$ -diffeomorphisms being a  $C^\infty$ -diffeomorphism, one can combine this transformation into a mixture to model a more complex bijection on  $[0, 1]$  (see Figure 4 in appendix). Because the inverse of mixtures of sigmoidal functions is not analytically defined, following Köhler et al. (2021) we compute the inverse of this coupling function numerically by the Newton-Raphson algorithm, and gradients through the inverse are obtained by the implicit function theorem (Blondel et al., 2021).

Figure 1 shows that using the smooth coupling function defined in (10) (right plot) it is possible to train the NF by penalizing its score through the score matching loss (6), which was not possible for an affine coupling function (middle plot).

## 4. Experiments

We tested our method using the combined loss (8) on two different tasks:

**Two-moons:** As a first illustration of the benefit of having access to the score to better constrain a distribution from a small number of samples in a non-conditional case, we consider a classical toy model called two-moons that consists a simple 2D bi-modal distribution. We are interested in learning the two-moons distribution  $p(x)$  from samples  $x \sim p(x)$  in two different cases: with simulations  $x$  only, and with simulations and score  $\nabla_x p(x)$ .

We use 3 coupling layers in our NF. The transformation parameters  $a, b, c$ , defined in (10), are learned using a neural network with 4 hidden layers of 128 units each and  $\sin$  activation functions. We trained our NF for six different training sample simulation budgets  $\{20, 50, 100, 200, 500, 1000\}$  and for each simulation budget we ran 10 realizations to compute the epistemic uncertainty. We evaluated the quality of the approximated distribution with the negative log-

likelihood.

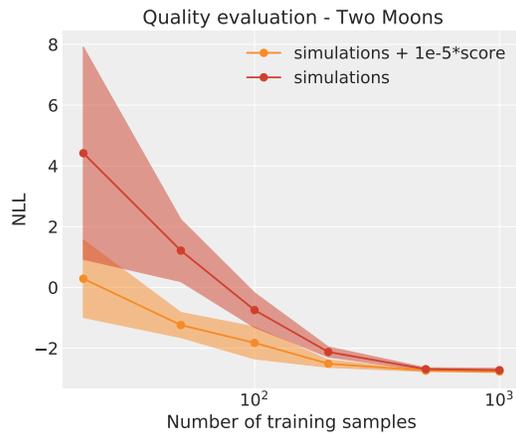


Figure 2. Two-moons inference: evolution of the quality of the distribution approximation with and without score information, measured by the negative log-likelihood, as a function of the size of the training set.

We find that in this simple case the gradients help to significantly accelerate the convergence of the density estimation. As shown on Figure 2, the approximated distribution obtained using simulation and score converge with only 200 simulations while 500 are needed for simulations alone. A comparison of the learned distributions as a function of the training set size is available in appendix in Figure 6.

**Lotka-Volterra:** We now consider the case of *conditional* density estimation with the classical Lotka-Volterra simulation model, which consists of a system of two non-linear differential equations. It describes the interaction between two different species, the preys  $X$  and the predators  $Y$ :

$$\begin{aligned} \frac{dX}{dt} &= \alpha X - \beta XY; \\ \frac{dY}{dt} &= -\gamma Y + \delta XY, \end{aligned} \quad (13)$$

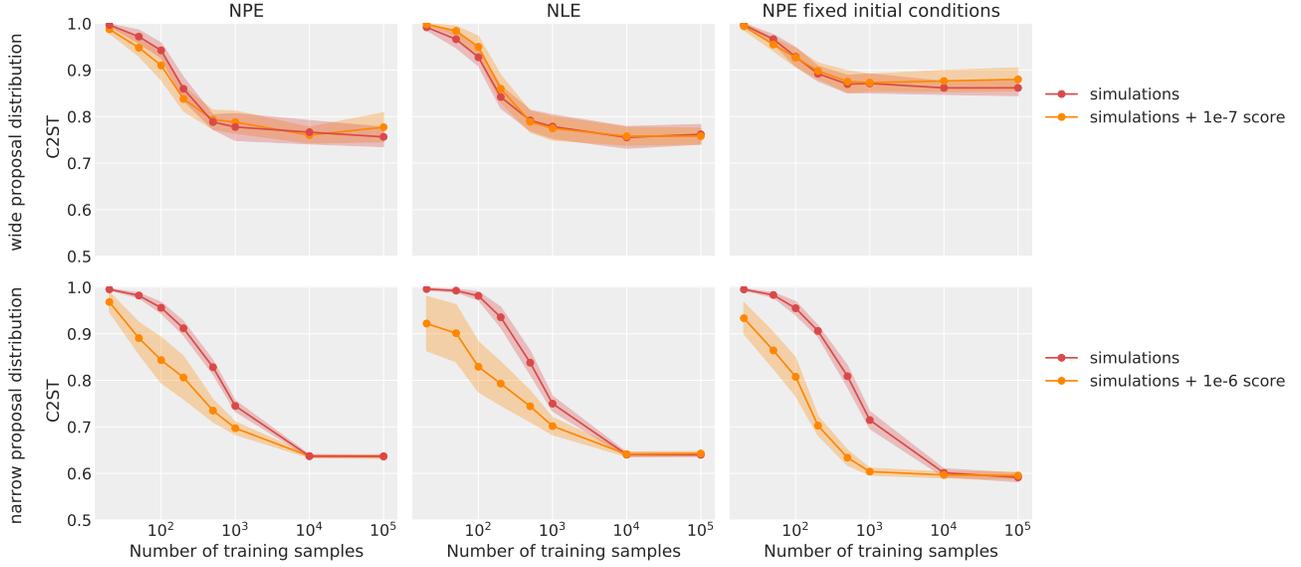


Figure 3. Lotka Volterra inference: quality of the posterior approximation (C2ST metric) as a function of the number of training samples used, for three methods with (orange) and without the score (red). Experiments from the top row have used a wide proposal distribution (16), while the one from the bottom row have used a narrow proposal distribution (17). First column: NPE. Second column: NLE. Last column: NPE for Lotka Volterra task with fixed initial conditions  $(x_0, Y_0) = (10, 1)$ .

with,  $\alpha, \beta, \gamma$  and  $\delta$  the inferred parameters  $\theta \in \mathbb{R}^4$ .

The simulations are 20 statistics that correspond to the number of each population  $X$  and  $Y$  at 10 different times. Hence our simulations correspond to  $x = (x_1, \dots, x_{10}, y_1, \dots, y_{10}) \in \mathbb{R}^{20}$ . We considered the initial conditions as our latent variables  $z$  with prior

$$z = \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} \sim \text{LogNormal}(\log(3), 0.5), \quad (14)$$

which lead to the following simulations

$$x|\theta = \begin{bmatrix} x_1, & \dots, & x_{10} \\ y_1, & \dots, & y_{10} \end{bmatrix} \sim \text{LogNormal}\left(\log \begin{bmatrix} X \\ Y \end{bmatrix}, 0.1\right). \quad (15)$$

We expect that the quality of the posterior approximation  $p(\theta|x)$  no longer depends only on capturing the shape of the distribution, but also on successfully conditioning it on observations. To illustrate these two factors, we design two experiments, one with a wide proposal distribution (compare to the posterior)

$$\begin{bmatrix} \alpha \\ \beta \\ \gamma \\ \delta \end{bmatrix} \sim \text{LogNormal}\left(\begin{bmatrix} -0.125 \\ -3 \\ -0.125 \\ -3 \end{bmatrix}, 0.5\right), \quad (16)$$

for which we expect errors from the conditioning to domi-

nate, and a narrow proposal distribution,

$$\begin{bmatrix} \alpha \\ \beta \\ \gamma \\ \delta \end{bmatrix} \sim \text{LogNormal}\left(\begin{bmatrix} -0.125 \\ -3 \\ -0.125 \\ -3 \end{bmatrix}, 0.03\right), \quad (17)$$

for which the main task is to focus on the posterior shape rather than finding its location.

As a first step, it is standard practice in SBI to compress  $x$  using an external neural network into a summary statistic of same dimensionality as the number of parameters. Moreover, this will allow us to make a fairer comparison between Neural Likelihood Estimation (NLE) and our NPE approach since in both cases the density estimation problem will be of same dimensionality and thus the same NF architecture can be used. In our case, we compressed  $x$  using a 1D CNN  $r_\phi$  trained externally. We trained the CNN under a NLL loss into a low dimensional  $4 - d$  summary statistics  $y$ , and used 4 smooth transformations in our NF. The transformation parameters  $a, b, c$ , defined in (10), were learned using a small NN (enough due to the external compressor mentioned above): 2 hidden layers of 128 units and *silu* activation functions (Elfving et al., 2017). To make our NDE conditional on  $y$ , this variable is used as an additional input to the coupling layers  $x$ . This yields the approximated distribution  $p_\phi(\theta|y) = r_\phi(x)$ .

We compared our method (NPE with score) with SCANDAL (NLE with score from Brehmer et al., 2020). We used the

same compressor and NF architecture and run a MCMC to get the approximated posterior from the learned likelihood  $p_\varphi(y = r_\phi(x)|\theta)$ .

For this task, we trained the NFs for eight different training sample simulations budget  $\{20, 50, 100, 200, 500, 10^3, 10^4, 10^5\}$  and for each simulation budget we ran 20 realizations to compute the epistemic uncertainty. We sample these simulations according to two different proposal distributions described above. We evaluated the quality of the approximated posterior with the Classifier 2-Sample Tests (C2ST) metric (Lueckmann et al., 2021). If the two samples come from the same distribution the classifier’s prediction is 0.5.

The two figures on the left in the first row of Figure 3 show that, in the case of a wide proposal distribution, adding the gradients in the inference process do not help to constrain the posterior. The score information does not help to constrain the conditioning of the density, i.e. to help constrain the overall location of the posterior mass. However, in the case of a narrow proposal distribution, the two figures on the left in the second row of the Figure 3, we find that score information significantly helps accelerating inference, for both NPE and NLE. This can be understood as in this case the location of the posterior mass is already well defined, and the score information contributes to constraining the shape of the posterior.

Finally, we highlight the impact of the stochastic score  $\nabla_\theta \log p(\theta|x, z)$  by repeating the same experiment with fixed initial conditions  $(X_0, Y_0) = (10, 1)$ . We found that the gradients help more without latent variables in the case of a narrow proposal distribution (see the last figures in the second row of Figure 3). This result is expected as the score in the fixed latent variable case is no longer stochastic, and can directly help to constrain the posterior distribution. However, fixed initial conditions make the posterior narrower and thus in the case of wide proposal distribution (see the last figures in the first row of Figure 3) the gradients help even less.

## 5. Discussion

To train the NF using simulations and score, we used a smooth NF from Köhler et al. (2021). Other types of differentiable  $C^\infty$  NF, such as ODE Flows Grathwohl et al. (2018), could be used in practice.

We demonstrated that having access to gradients from the simulation is only beneficial for the density shape estimation. This points to a high interest of having access to score information to refine rough estimates, and thus to sequential density estimation methods. We expect these tools to become increasingly useful in astrophysics with the advent of automatically differentiable physical simulators like the FlowPM cosmological simulation code (Modi et al., 2020).

## References

- Blondel, M., Berthet, Q., Cuturi, M., Frostig, R., Hoyer, S., Llinares-López, F., Pedregosa, F., and Vert, J.-P. Efficient and modular implicit differentiation, 2021. URL <https://arxiv.org/abs/2105.15183>.
- Blum, M. G. B. and François, O. Non-linear regression models for approximate bayesian computation. *Statistics and Computing*, 20(1):63–73, mar 2009. doi: 10.1007/s11222-009-9116-0. URL <https://doi.org/10.10072Fs11222-009-9116-0>.
- Brehmer, J., Louppe, G., Pavez, J., and Cranmer, K. Mining gold from implicit models to improve likelihood-free inference. *Proceedings of the National Academy of Sciences*, 117(10):5242–5249, feb 2020. doi: 10.1073/pnas.1915980117. URL <https://doi.org/10.1073%2Fpnas.1915980117>.
- Cranmer, K., Pavez, J., and Louppe, G. Approximating likelihood ratios with calibrated discriminative classifiers, 2015. URL <https://arxiv.org/abs/1506.02169>.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation, 2014. URL <https://arxiv.org/abs/1410.8516>.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp, 2016. URL <https://arxiv.org/abs/1605.08803>.
- Durkan, C., Murray, I., and Papamakarios, G. On contrastive learning for likelihood-free inference, 2020. URL <https://arxiv.org/abs/2002.03712>.
- Elfving, S., Uchibe, E., and Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning, 2017.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models, 2018. URL <https://arxiv.org/abs/1810.01367>.
- Greenberg, D. S., Nonnenmacher, M., and Macke, J. H. Automatic posterior transformation for likelihood-free inference, 2019.
- Izbicki, R., Lee, A. B., and Schafer, C. M. High-dimensional density ratio estimation with extensions to approximate likelihood computation. 2014. doi: 10.48550/ARXIV.1404.7063. URL <https://arxiv.org/abs/1404.7063>.
- Köhler, J., Krämer, A., and Noé, F. Smooth normalizing flows, 2021. URL <https://arxiv.org/abs/2110.00351>.

- 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, 2017.
- Lueckmann, J.-M., Bassetto, G., Karaletsos, T., and Macke, J. H. Likelihood-free inference with emulator networks. 2018. doi: 10.48550/ARXIV.1805.09294. URL <https://arxiv.org/abs/1805.09294>.
- Lueckmann, J.-M., Boelts, J., Greenberg, D. S., Gonçalves, P. J., and Macke, J. H. Benchmarking simulation-based inference, 2021. URL <https://arxiv.org/abs/2101.04653>.
- Modi, C., Lanusse, F., and Seljak, U. Flowpm: Distributed tensorflow implementation of the fastpm cosmological n-body solver, 2020. URL <https://arxiv.org/abs/2010.11847>.
- Papamakarios, G. and Murray, I. Fast  $\epsilon$ -free inference of simulation models with bayesian conditional density estimation, 2018.
- Papamakarios, G., Sterratt, D. C., and Murray, I. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows, 2018. URL <https://arxiv.org/abs/1805.07226>.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows, 2015. URL <https://arxiv.org/abs/1505.05770>.
- Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. Likelihood-free inference by ratio estimation, 2016. URL <https://arxiv.org/abs/1611.10242>.
- Wood, S. N. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, August 2010. ISSN 1476-4687. doi: 10.1038/nature09319. URL <https://doi.org/10.1038/nature09319>.

A. Additional plots

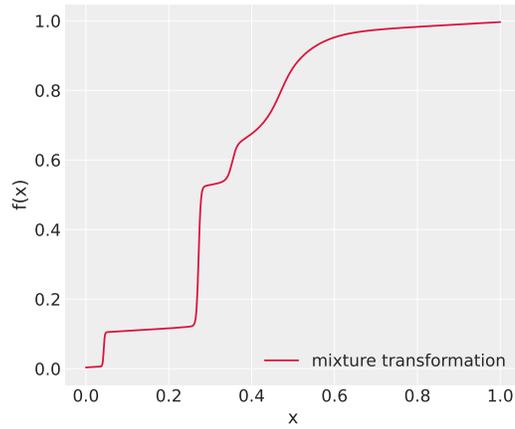


Figure 4. Mixture of 5 smooth transformations used to create a smooth Normalizing Flow.

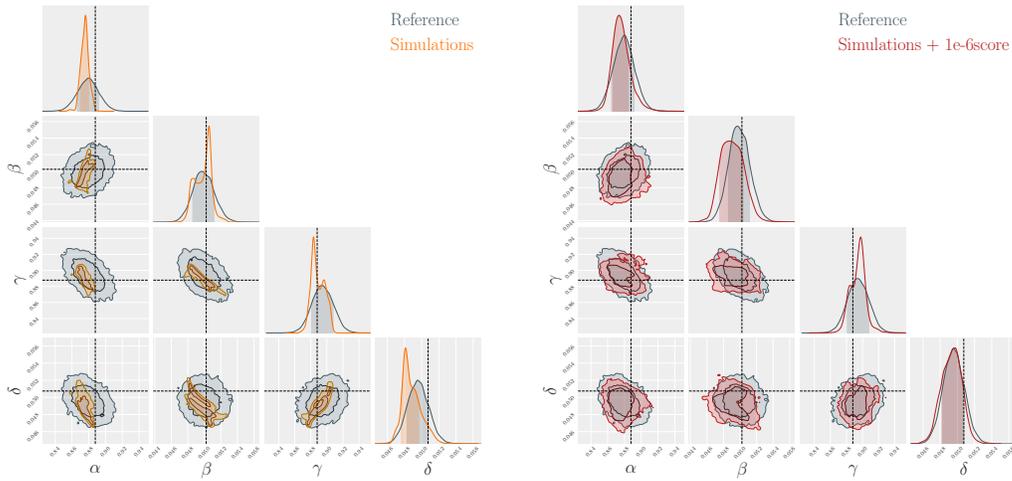


Figure 5. Lotka Volterra posterior comparison between the reference posterior (blue), the posterior obtained using only simulations (orange on the left) and the one obtained using both simulations and gradients (red on the right). These posteriors are approximated using NPE method with only 50 simulations in the case of a narrow proposal distribution.

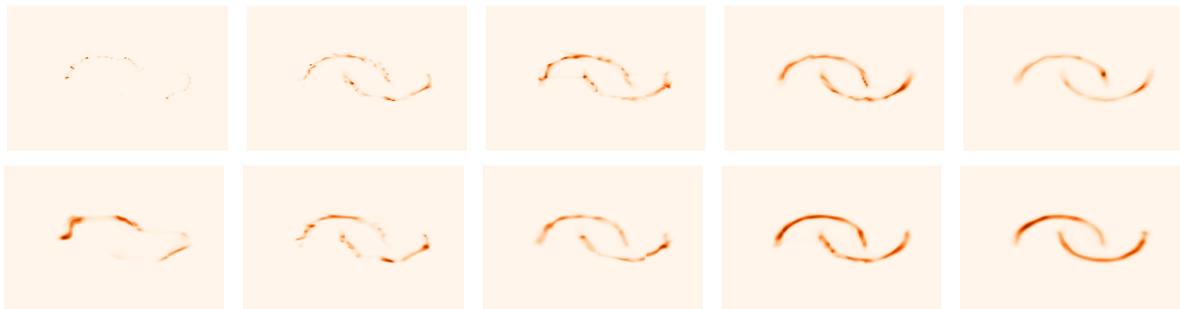


Figure 6. Two Moons distribution approximated using, from left to right, 20, 100, 200, 500, 1000 simulations. The first row correspond to the distribution learned using simulations only and the second row learned on simulations and score.