

---

# Deformations of Boltzmann Distributions

---

**Bálint Máté**  
University of Geneva  
balint.mate@unige.ch

**François Fleuret**  
University of Geneva  
francois.fleuret@unige.ch

## Abstract

Consider a one-parameter family of Boltzmann distributions  $p_t(x) = \frac{1}{Z_t} e^{-S_t(x)}$ . This work studies the problem of sampling from  $p_{t_0}$  by first sampling from  $p_{t_1}$  and then applying a transformation  $\Psi_{t_1}^{t_0}$  so that the transformed samples follow  $p_{t_0}$ . We derive an equation relating  $\Psi$  and the corresponding family of unnormalized log-likelihoods  $S_t$ . The utility of this idea is demonstrated on the  $\phi^4$  lattice field theory by extending its defining action  $S_0$  to a family of actions  $S_t$  and finding a  $\tau$  such that normalizing flows perform better at learning the Boltzmann distribution  $p_\tau$  than at learning  $p_0$ .

## 1 Introduction

Sampling from unnormalized densities has been studied by many due to its relevance for the sciences [1–11]. The problem can be summarized as follows. Given an unnormalized log-density  $S : \mathbb{R}^n \rightarrow \mathbb{R}$  can we efficiently generate samples from the probability density  $p(x) = \frac{1}{Z} e^{-S(x)}$ ? In particular, there are no samples given, all we have is the ability to evaluate  $S$  for any sample candidate. A popular technique for attacking this problem is to use a normalizing flow to parametrise a distribution  $q_\theta$  and optimize the parameters  $\theta$  to minimize the reverse KL divergence

$$KL[q_\theta, p] = \mathbb{E}_{x \sim q_\theta} (\log q_\theta(x) - \log p(x)) = \mathbb{E}_{x \sim q_\theta} (\log q_\theta(x) + S(x)) + Z. \quad (1)$$

As a motivating example for this paper, let  $S$  be the defining action of the lattice  $\phi^4$  theory (See §4 for details), and consider the family of distributions  $p_\beta(x) \propto e^{-\beta S(x)}$  parametrized by  $\beta \in \mathbb{R}^+$ . In terms of statistical physics  $\beta$  corresponds to the inverse temperature and controls how ordered the given system is. As seen in Fig. 1 the performance of a normalizing flow is sensitive to the parameter  $\beta$ . Continuous normalizing flows converge faster at higher temperatures. Somewhat surprisingly, the RealNVP architecture converges faster both at  $\beta = .1$  and  $\beta = 10$  than at  $\beta = 1$ . A possible explanation is that both the smoother distribution ( $\beta = 0.1$ ) and the more localized ( $\beta = 10$ ) is easier to learn than a combination of these characteristics at  $\beta = 1$ .

Motivated by this observation, this paper studies the following problem. Suppose we are given a family of actions  $S_t(x)$ , parametrised by  $t$ , inducing a family of distributions,  $p_t(x) \propto e^{-S_t(x)}$ . We are interested in how these distributions are related for different values of  $t$ . From a practical viewpoint, the goal is to sample from  $p_{t_2}$  by first sampling from  $p_{t_1}$  and making the samples flow along some vector field  $V_t$ . We will refer to this  $V_t$  as the deformation field or transport field.

The contributions of this paper can be summarized as:

- In §3 we derive a PDE, the deformation equation, that translates between infinitesimal deformations of actions and the infinitesimal deformations of the Boltzmann distributions they induce.
- In §4 we put the theory into practice in the case of a simple deformation of the lattice  $\phi^4$  theory and show that it improves the performance of normalizing flows.

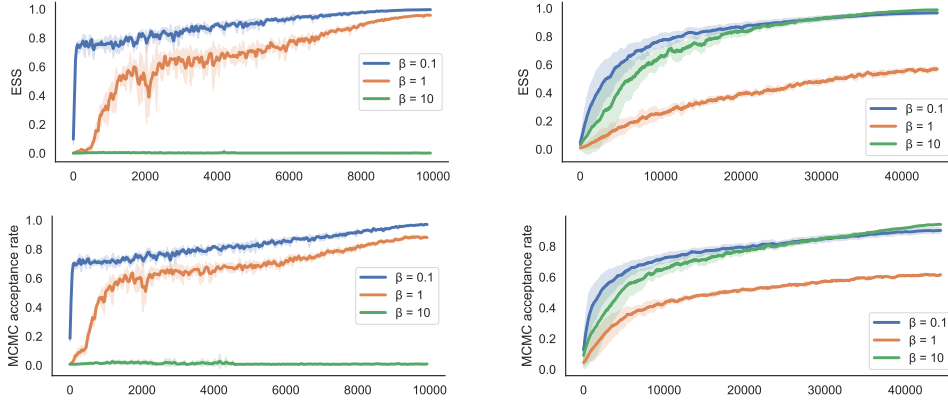


Figure 1: Sensitivity of the training to the inverse temperature  $\beta$  on a  $12 \times 12$  lattice. We trained the continuous normalizing flow of Gerdes et al. [7] (left) and a RealNVP (right). In both cases we used the action of the  $\phi^4$  theory with different values of  $\beta \in \{0.1, 1, 10\}$  and with  $m^2$  and  $\lambda$  values same as in the work of Gerdes et al. [7]. The  $x$ -axis represents the number of training steps, while the  $y$ -axis represents the performance metrics. Mean and standard deviation over 5 runs are shown.




Figure 2: Deformation of a two-dimensional standard Gaussian into a family of isotropic Gaussians with different variances. Samples from a two-dimensional standard Gaussian (left). The deformation field (center). The samples in the left plot flowing along the deformation field from 1 to  $t$  follow a Gaussian centered at the origin with covariance  $1/t$  (right).

## 2 Background

**Change of variables** Let  $p_0(z)$  be a probability density on  $\mathbb{R}^n$  and  $\Psi : \mathbb{R}^n \rightarrow \mathbb{R}^n$  a diffeomorphism. Pushing the the density forward along  $\Psi$  induces a new probability density  $p$  implicitly defined by

$$\log p_0(z) = \log p(\Psi z) + \log |\det J_\Psi(z)|. \quad (2)$$

The term  $\log |\det J_\Psi(z)|$  measures how much the function  $\Psi$  expands volume locally at  $z$ .

**Continous change of variables** Suppose  $V_t$  is a time-dependent vector field. Let  $\Psi_\tau$  denote the diffeomorphism of following the trajectories of  $V_t$  from 0 to  $\tau$ . This family of diffeomorphisms generates a one-parameter family of densities  $p_\tau$ . The amount of volume expansion a particle experiences along this trajectory of  $\Psi_\tau$  is  $\int_0^\tau \text{div } V_t(\Psi_t z) dt$ . The log-likelihoods are then related by

$$\log p_0(z) = \log p_\tau(\Psi_\tau z) + \int_0^\tau \text{div } V_t(\Psi_t z) dt. \quad (3)$$

**Normalizing flows** Normalizing flows [12–15] (continuous normalizing flows [16], respectively) parametrize a subset of the space of all distributions on  $\mathbb{R}^n$ . They do this by first fixing a base density  $p_0$  and using a neural network that parametrizes the transformation  $\Psi$  (the vector field  $V_t$ , respectively). The change of variables formula in Eq. 2 (Eq. 3, respectively) is then applied to compute the distribution induced by  $\Psi$  ( $V_t$ , respectively).

**Performance metrics** To evaluate the performance of normalizing flows, we employ the following two metrics. First, let  $q_\theta$  be the distribution parametrized by a normalizing flow and  $p \propto e^{-S}$  the target distribution. Given a batch of samples  $x_i$ , the effective sample size can be computed as

$$ESS = \frac{\left(\frac{1}{N} \sum_i p[x_i]/q_\theta[x_i]\right)^2}{\frac{1}{N} \sum_i (p[x_i]/q_\theta[x_i])^2}. \quad (4)$$

We now think of the samples as being generated sequentially in a MCMC manner. We always accept the first sample  $x_0$ , and each of the following samples are accepted with probability

$$p_{accept}(x^i|x^{i-1}) = \min\left(1, \frac{q_\theta[x^{i-1}]}{p[x^{i-1}]} \frac{p[x^i]}{q_\theta[x^i]}\right). \quad (5)$$

Note that all ratios  $p[x]/q_\theta[x]$  are invariant to change of coordinates because both the numerator and denominator get multiplied with determinant of the same Jacobian. This in turn implies that both the ESS and the MCMC acceptance probability are invariant under a change of coordinates.

**Boltzmann distributions** Let  $S : \mathbb{R}^n \rightarrow \mathbb{R}$  be a function with a finite normalizing constant  $Z = \int e^{-S(x)} d^n x$ . Physically speaking,  $S$  might be the action of some physical theory or the energy associated to configurations of some system. Either way,  $S$  induces a Boltzmann distribution over the configurations  $x \in \mathbb{R}^n$

$$p(x) = \frac{1}{Z} e^{-S(x)}. \quad (6)$$

Conversely, given a probability density function  $p : \mathbb{R}^n \rightarrow \mathbb{R}_{+,0}$  the corresponding action can be recovered up to a constant  $S = -\log p - \log Z$  and Eqs. 2 and 3 can be interpreted as equations relating (the actions and normalizing constants of) two physical theories over the same configuration space.

### 3 The deformation equation

In what follows, we will use the symbols  $\partial_t f$  and  $\nabla f$  to denote the derivatives with respect to the time parameter and the gradient over the remaining  $n$ , spatial components of a time-dependent scalar valued function  $f : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ .

The Boltzmann distribution of Eq. 6 can be equivalently defined by setting the ratios of likelihoods between all  $x$  and  $y$  to

$$\frac{p(x)}{p(y)} = e^{S(y)-S(x)}. \quad (7)$$

Our aim is then to construct the time-dependent vector field  $V_t$  such that flowing along  $V_t$  from  $t_1$  to  $t_2$  transforms between the two densities

$$p_{t_1}(x) = p_{t_2}(\Psi_{t_1}^{t_2} x) \mathcal{L}_{\Psi_{t_1}^{t_2}}(x), \quad (8)$$

where  $\Psi_{t_1}^{t_2}$  denotes the operation of "flowing along  $V_t$  from  $t_1$  to  $t_2$ " and  $\mathcal{L}_{\Psi_{t_1}^{t_2}}$  is the likelihood contribution of  $\Psi_{t_1}^{t_2}$ . In terms of Eq. 7, we are looking for  $\Psi$  such that for any pair of  $x, y$  we have

$$\frac{p_{t_1}(x)}{p_{t_1}(y)} = \frac{p_{t_2}(\Psi x) \mathcal{L}_\Psi(x)}{p_{t_2}(\Psi y) \mathcal{L}_\Psi(y)}, \quad (9)$$

where we dropped the indices from  $\Psi$  for readability. For Boltzmann distributions, this expression is equivalent to

$$e^{S_{t_1}(y)-S_{t_1}(x)} = e^{S_{t_2}(\Psi y)-S_{t_2}(\Psi x)} \frac{\mathcal{L}_\Psi(x)}{\mathcal{L}_\Psi(y)}. \quad (10)$$

Taking the logarithm and rearranging

$$S_{t_1}(x) - S_{t_2}(\Psi x) + \log \mathcal{L}_\Psi(x) = S_{t_1}(y) - S_{t_2}(\Psi y) + \log \mathcal{L}_\Psi(y). \quad (11)$$

This is satisfied if and only if the expression (with reinserted indices on  $\Psi$ )

$$f_{t_1}^{t_2}(x) = S_{t_1}(x) - S_{t_2}(\Psi_{t_1}^{t_2} x) + \log \mathcal{L}_{\Psi_{t_1}^{t_2}}(x), \quad (12)$$

is independent of  $x$ . Which happens if and only if  $\nabla f_{t_1}^{t_2}(x)$  vanishes for all  $x, t_1, t_2$ . We now construct the vector field  $V_t$ , the infinitesimal generator of  $\Psi$ . Expanding  $\nabla f_{t_1}^{t_2}$  in  $t_2$  at  $t_1$  up to first order yields

$$\nabla f_{t_1}^{t_2}(x) = \underbrace{\nabla f_{t_1}^{t_2}(x)}_{\Big|_{t_2=t_1} \quad 0} + \underbrace{\partial_{t_2} \nabla f_{t_1}^{t_2}(x)}_{\nabla \partial_{t_2}} \Big|_{t_2=t_1} \delta t + \mathcal{O}(\delta t^2) \quad (13)$$

$$= \nabla [-\partial_t S_{t_1}(x) - S_{t_1}(\partial_{t_2} \Psi_{t_1}^{t_2} x) + \text{div } V(x)] \Big|_{t_2=t_1} \delta t + \mathcal{O}(\delta t^2) \quad (14)$$

$$= \nabla [-\partial_t S_{t_1}(x) - \langle \nabla S_{t_1}, V_{t_1}(x) \rangle + \text{div } V(x)] \delta t + \mathcal{O}(\delta t^2). \quad (15)$$

For our purposes  $\nabla f_{t_1}^{t_2}$  must vanish everywhere, therefore  $V_t$  must be a solution of

$$\nabla [\partial_t S_t + \langle \nabla S_t, V_t \rangle - \text{div } V_t] = 0 \Leftrightarrow \partial_t S_t + \langle \nabla S_t, V_t \rangle - \text{div } V_t \approx 0, \quad (16)$$

where the symbol  $\approx$  means that the two sides only agree up to a spatially constant function. A closely related equation is derived by Pfau and Rezende [17, Eq. 6]. Eq. 16 can also be obtained by taking the gradient of their Eq. 6, removing the  $t$ -dependent normalizing constant  $Z_t$ .

Given an infinitesimal deformation  $\partial_t S$  of the action  $S$ , the vector field  $V_t$  generating  $\partial_t S$  is obtained by solving Eq. 16 for  $V_t$ . Conversely, given a vector field  $V_t$ , the infinitesimal deformation it induces on  $S$  is found by solving the same equation for  $S_t$ .

### Sanity check: 1D Gaussians

In general, given a family of actions  $S_t$ , it is difficult to solve Eq. 16 for  $V_t$ . In the following toy-example we explicitly find a particular solution. We consider one-dimensional centered Gaussians that are defined by the action

$$S_t(x) = t \frac{1}{2} x^2 \quad \Rightarrow \quad p_t(x) = \frac{1}{Z_t} e^{-\frac{t}{2} x^2} = \mathcal{N}(\mu = 0, \sigma^2 = 1/t). \quad (17)$$

Changing  $t$  to  $t + \epsilon$  changes the distribution from  $\mathcal{N}(0, 1/t)$  to  $\mathcal{N}(0, 1/(t + \epsilon))$ . To map samples from the former to the latter, we can multiply with the ratios of the standard deviations, i.e. apply the map  $x \mapsto x \sqrt{\frac{t}{t + \epsilon}}$ . The transport field  $V_t$  is then given by the derivative of this map with respect to  $\epsilon$ ,

$$V_t = \frac{d}{d\epsilon} \left[ x \sqrt{\frac{t}{t + \epsilon}} \right] \Big|_{\epsilon=0} = -\frac{x}{2t}. \quad (18)$$

This choice for  $V_t$  also solves the deformation equation of the previous section. Indeed, plugging  $S_t = t \frac{1}{2} x^2$  and  $V_t = -\frac{x}{2t}$  into Eq. 16,

$$\partial_x [\partial_t S_t + \langle \nabla S_t, V_t \rangle - \text{div } V_t] = \partial_x \left[ \frac{1}{2} x^2 - tx \frac{x}{2t} + \frac{1}{2t} \right] = 0. \quad (19)$$

This example generalizes to higher dimensional Gaussians, a two dimensional example can be seen in Fig. 2.

## 4 The $\phi^4$ theory on a 2D lattice

The  $\phi^4$  theory is defined by the action  $S_0[\phi] = \int \|\nabla\phi\|^2 + m^2\phi^2 + \lambda\phi^4 dx$ . After discretizing to a finite lattice  $L$ , the action becomes

$$\sum_{x,y \in L} \phi(x)\Delta_{x,y}^L\phi(y) + \sum_{x \in L} m^2\phi^2(x) + \sum_{x \in L} \lambda_0\phi^4(x) = \phi^T \underbrace{(\Delta^L + m^2)}_{K_0} \phi + \lambda_0 \|\phi\|_4^4, \quad (20)$$

where  $\Delta^L$  is the Laplacian of the lattice  $L$ . We now deform with one of the simplest non-trivial vector fields,  $V_t(\phi) = V(\phi) = \phi$ . As a function,  $V$  is just the identity, and the dynamics it induces is a radial expansion centered at the origin. Plugging it into Eq. 16,

$$\partial_t S_t = \text{div } V - \langle \nabla S_t, V \rangle = \dim \phi - \langle 2K_t\phi + 4\lambda_t\phi^3, \phi \rangle \approx -2\phi^T K_t\phi - 4\lambda_t \|\phi\|_4^4, \quad (21)$$

where we assumed that  $S_t$  is of the form  $S_t = \phi^T K_t\phi + \lambda_t \|\phi\|_4^4$ . Indeed, setting  $K_t = e^{-2t}K_0$  and  $\lambda_t = e^{-4t}\lambda_0$  yields a solution,

$$S_t[\phi] = e^{-2t}(\phi^T K_0\phi) + e^{-4t}(\lambda_0 \|\phi\|_4^4). \quad (22)$$

An interesting point to note here is the similarity between the family of actions  $S_t$  given by Eq. 22 and the family  $S_\beta = \beta S_0$  that motivated this paper. The family  $S_\beta$  deforms  $S_0$  by scaling the whole action with the same multiplicative constant, while the family  $S_t$  deforms  $S_0$  by scaling quartic terms quadratically faster than the quadratic terms. Fig. 3 shows how the deformation in Eq. 22 with a positive value of  $t$  improves the performance of normalizing flows.

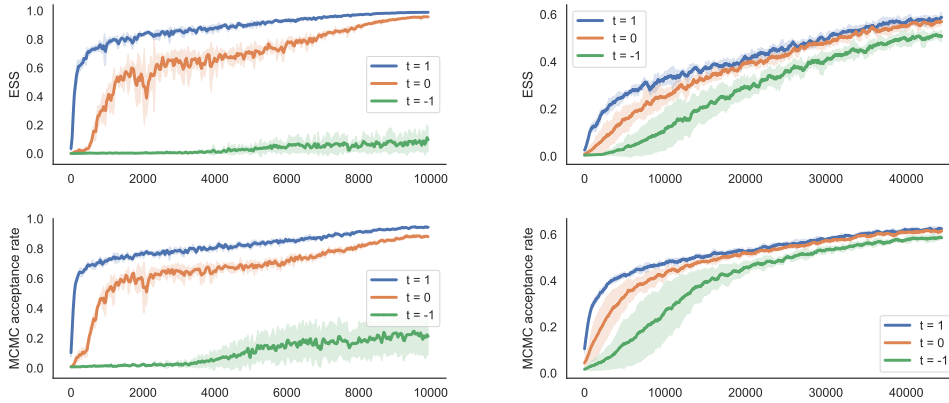


Figure 3: Sensitivity of the training to the deformation parameter  $t$  on a  $12 \times 12$  lattice. We trained the continuous normalizing flow of Gerdes et al. [7] (left) and a RealNVP (right). In both cases we used the action given in Eq. 22 with different values of  $t \in \{-1, 0, 1\}$  and with  $m^2$  and  $\lambda$  values same as in the work of Gerdes et al. [7]. The  $x$ -axis represents the number of training steps, while the  $y$ -axis represents the performance metrics. Mean and standard deviation over 5 runs are shown.

## 5 Conclusion and Impact Statement

**Conclusion and limitations** We propose a technique that deforms one Boltzmann distribution into another that is easier to model with normalizing flows. We experimentally support this in the case of the  $\phi^4$  lattice field theory. Our experiments indicate that the technique leads to improved ESS and acceptance rates. It is important to note however that good ESS and acceptance rate values does not mean that  $q_\theta$  is perfectly matching  $p$ . For example, if there are several modes of  $p$ ,  $q_\theta$  could be missing some of the modes while still obtaining good ESS and acceptance rate values.

**Impact Statement** Beyond the application in lattice field theory that we presented in this paper, the technique of deforming distributions naturally lends itself to being used for modelling other physical and chemical systems. The potential impact of the technique introduced in this paper is to allow the above tasks to be solved more efficiently. In particular, the authors do not see any ethical concerns nor a direct potential of negative impact to society.

## 6 Acknowledgement

The authors acknowledge support from the Swiss National Science Foundation under grant number CRSII5\_193716 - "Robust Deep Density Models for High-Energy Particle Physics and Solar Flare Analysis (RODEM)". We further thank Samuel Klein and Jonas Köhler for discussions.

## References

- [1] Denis Boyda, Gurtej Kanwar, Sébastien Racanière, Danilo Jimenez Rezende, Michael S. Albergo, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan. Sampling using  $SU(n)$  gauge equivariant flows. *Phys. Rev. D*, 103:074504, Apr 2021. doi: 10.1103/PhysRevD.103.074504. URL <https://link.aps.org/doi/10.1103/PhysRevD.103.074504>.
- [2] Michael S. Albergo, Denis Boyda, Daniel C. Hackett, Gurtej Kanwar, Kyle Cranmer, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan. Introduction to normalizing flows for lattice field theory, 2021. URL <https://arxiv.org/abs/2101.08176>.
- [3] Michael S. Albergo, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Julian M. Urban, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan. Flow-based sampling for fermionic lattice field theories. *Physical Review D*, 104(11), dec 2021. doi: 10.1103/physrevd.104.114507. URL <https://doi.org/10.1103/physrevd.104.114507>.
- [4] Michael S. Albergo, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Fernando Romero-López, Phiala E. Shanahan, and Julian M. Urban. Flow-based sampling in the lattice schwinger model at criticality, 2022. URL <https://arxiv.org/abs/2202.11712>.
- [5] Ryan Abbott, Michael S. Albergo, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Fernando Romero-López, Phiala E. Shanahan, Betsy Tian, and Julian M. Urban. Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions, 2022. URL <https://arxiv.org/abs/2207.08945>.
- [6] Pim de Haan, Corrado Rainone, Miranda C. N. Cheng, and Roberto Bondesan. Scaling up machine learning for quantum field theory with equivariant continuous flows, 2021. URL <https://arxiv.org/abs/2110.02673>.
- [7] Mathis Gerdes, Pim de Haan, Corrado Rainone, Roberto Bondesan, and Miranda C. N. Cheng. Learning lattice quantum field theories with equivariant continuous flows, 2022. URL <https://arxiv.org/abs/2207.00283>.
- [8] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators – sampling equilibrium states of many-body systems with deep learning, 2018. URL <https://arxiv.org/abs/1812.01729>.
- [9] Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: Exact likelihood generative learning for symmetric densities, 2020. URL <https://arxiv.org/abs/2006.02425>.
- [10] Kim A. Nicoli, Shinichi Nakajima, Nils Strodthoff, Wojciech Samek, Klaus-Robert Müller, and Pan Kessel. Asymptotically unbiased estimation of physical observables with neural samplers. *Physical Review E*, 101(2), feb 2020. doi: 10.1103/physreve.101.023304. URL <https://doi.org/10.1103/physreve.101.023304>.
- [11] Kim A Nicoli, Christopher J Anders, Lena Funcke, Tobias Hartung, Karl Jansen, Pan Kessel, Shinichi Nakajima, and Paolo Stornati. Estimation of thermodynamic observables in lattice field theories with deep generative models. *Physical review letters*, 126(3):032001, 2021.
- [12] Esteban G Tabak and Cristina V Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- [13] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

- [14] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.
- [15] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows, 2019.
- [16] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations, 2018. URL <https://arxiv.org/abs/1806.07366>.
- [17] David Pfau and Danilo Rezende. Integrable nonparametric flows. *arXiv preprint arXiv:2012.02035*, 2020.
- [18] Ricky T. Q. Chen. torchdiffeq, 2018. URL <https://github.com/rtqichen/torchdiffeq>.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [N/A]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Experimental details

All of our experiments were implemented in PyTorch. We trained all the models with a sample size of 128 at each training step. The learning rate was initialized to  $3 \times 10^{-3}$  and annealed to 0 following a cosine schedule.

**RealNVP** The RealNVP architectures use an alternating checkerboard masking and a 3-layer convolutional network in each layer with kernel size 3 to compute the parameters of the affine transformations. The overall architecture contains 24 such layers.

**Continuous NF** The experiments involving continuous normalizing flows used the architecture proposed by Gerdes et al. [7] with  $F = 16$  field basis functions and  $D = 9$  time kernels, and coupling dimensions  $F' = 8, D' = 8$ . The 4th order Runge-Kutta solver of the torchdiffeq package [18] with a step size of 0.02 was used.