

---

# Block Neural Autoregressive Flow

---

**Nicola De Cao**  
University of Edinburgh  
University of Amsterdam  
nicola.decao@uva.nl

**Wilker Aziz**  
University of Amsterdam  
w.aziz@uva.nl

**Ivan Titov**  
University of Edinburgh  
University of Amsterdam  
ititov@inf.ed.ac.uk

## Abstract

Normalising flows (NFs) map two density functions via a differentiable bijection whose Jacobian determinant can be computed efficiently. Recently, as an alternative to hand-crafted bijections, [Huang et al. \(2018\)](#) proposed neural autoregressive flow (NAF) which is a universal approximator for density functions. Their flow is a neural network (NN) whose parameters are predicted by another NN. The latter grows quadratically with the size of the former and thus an efficient technique for parametrization is needed. We propose *block neural autoregressive flow* (B-NAF), a much more compact universal approximator of density functions, where we model a bijection directly using a single feed-forward network. Invertibility is ensured by carefully designing each affine transformation with block matrices that make the flow autoregressive and (strictly) monotone. We compare B-NAF to NAF and other established flows on density estimation and approximate inference for latent variable models. Our proposed flow is competitive across datasets while using orders of magnitude fewer parameters.

## 1 INTRODUCTION

Normalizing flows (NFs) map two probability density functions via an invertible transformation with tractable Jacobian ([Tabak et al., 2010](#)). They have been employed in contexts where we need to model a complex density while maintaining efficient sampling and/or density assessments. In density estimation, for example, NFs are used to map observations from a complex (and unknown) data distribution to samples from a simple base distri-

bution ([Rippel and Adams, 2013](#)). In variational inference, NFs map from a simple fixed random source (e.g. a standard Gaussian) to a complex posterior approximation while allowing for reparametrized gradient estimates and density assessments ([Rezende and Mohamed, 2015](#)).

Much of the research in NFs focuses on designing expressive transformations while satisfying practical constraints. In particular, autoregressive flows (AFs) decompose a joint distribution over  $y \in \mathbb{R}^d$  into a product of  $d$  univariate conditionals. A transformation  $y = f(x)$ , that realizes such a decomposition, has a lower triangular Jacobian with the determinant (necessary for application of the change of variables theorem for densities) computable in  $O(d)$ -time. [Kingma et al. \(2016\)](#) proposed *inverse autoregressive flows* (IAFs), an AF based on transforming each conditional by a composition of a finite number of trivially invertible affine transformations.

Recently, [Huang et al. \(2018\)](#) introduced *neural autoregressive flows* (NAFs). They replace IAF’s transformation by a learned bijection realized as a strictly monotonic neural network. Notably, they prove that their method is a universal approximator of real and continuous distributions. Though, whereas parametrizing affine transformations in an IAF requires predicting  $d$  pairs of scalars per step of the flow, parametrizing a NAF requires predicting all the parameters of a feed-forward *transformer* network. The *conditioner* network which parametrizes the transformer grows quadratically with the width of the transformer network, thus efficient parametrization techniques are necessary. A NAF is an instance of a hyper-network ([Ha et al., 2017](#)).

**Contribution** We propose *block neural autoregressive flows* (B-NAFs),<sup>1</sup> which are AFs based on a novel transformer network which transforms conditionals directly, i.e. without the need for a conditioner network. To do that we exploit the fact that invertibility only requires

---

<sup>1</sup><https://github.com/nicola-decao/BNAF>

$\partial y_i / \partial x_i > 0$ , and therefore, careful design of a feed-forward network can ensure that the transformation is both autoregressive (with unconstrained manipulation of  $x_{<i}$ ) and strictly monotone (with positive  $\partial y_i / \partial x_i$ ). We do so by organizing the weight matrices of dense layers in block matrices that independently transform subsets of the variables and constrain these blocks to guarantee that  $\partial y_i / \partial x_j = 0$  for  $j > i$  and that  $\partial y_i / \partial x_i > 0$ . Our B-NAFs are much more compact than NAFs while remaining universal approximators of density functions. We evaluate them both on density estimation and variational inference showing performance on par with state-of-the-art NFs, including NAFs, IAFs, and Sylvester flows (van den Berg et al., 2018), while using orders of magnitude fewer parameters.

## 2 BACKGROUND

In this section, we provide an introduction to normalizing flows and their applications (§ 2.1). Then, we motivate autoregressive flows in § 2.2 and present the necessary background for our contributions (§ 2.3).

### 2.1 NORMALIZING FLOW

A (finite) normalizing flow is a bijective function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  between two continuous random variables  $X \in \mathcal{X} \subseteq \mathbb{R}^d$  and  $Y \in \mathcal{Y} \subseteq \mathbb{R}^d$  (Tabak et al., 2010). The change of variables theorem expresses a relation between the probability density functions  $p_Y(y)$  and  $p_X(x)$ :

$$p_Y(y) = p_X(x) |\det \mathbf{J}_{f(x)}|^{-1}, \quad (1)$$

where  $y = f(x)$ , and  $|\det \mathbf{J}_{f(x)}|$  is the absolute value of the determinant of the Jacobian of  $f$  evaluated at  $x$ . The Jacobian matrix is defined as  $(\mathbf{J}_{f(x)})_{ij} = \partial f(x)_i / \partial x_j$ . The determinant quantifies how  $f$  locally expands or contracts regions of  $\mathcal{X}$ . Note that a composition of invertible functions remain invertible, thus a composition of NFs is itself a normalizing flow.

**Density estimation** In parametric density estimation, NFs map draws from complex distributions to draws from simple ones allowing assessments of a complex likelihood. Effectively, observations  $x \sim p_{\text{data}}$  are modeled as draws from an NF  $p_{X|\theta}$  whose parameters  $\theta$  are chosen to minimize the Kullback-Leibler divergence  $\mathbb{KL}(p_{\text{data}} \| p_{X|\theta})$  between the model  $p_{X|\theta}$  and  $p_{\text{data}}$ :

$$\mathbb{H}(p_{\text{data}}) - \mathbb{E}_{p_{\text{data}}} [\log p_Y(f(x)) |\det \mathbf{J}_{f_\theta(x)}|], \quad (2)$$

where  $\mathbb{H}$  indicates the entropy. Minimizing such KL is equivalent to maximizing the log-likelihood of observations (see Appendix A for details of such derivation).

**Variational inference** In variational inference for deep generative models, NFs map draws from a simple density  $q_X$ , e.g.  $\mathcal{N}(0, I)$ , to draws from a complex (multimodal) density  $q_{Y|\theta}(y)$ . The parameters  $\theta$  of the flow are estimated to minimize the KL divergence  $\mathbb{KL}(q_{Y|\theta} \| p_Y)$  between the model  $q_{Y|\theta}$  and the true posterior  $p_Y$ :

$$\mathbb{E}_{q_X(x)} \left[ \log \frac{q_X(x) |\det \mathbf{J}_{f_\theta(x)}|^{-1}}{p_Y(f_\theta(x))} \right], \quad (3)$$

and note that this enables backpropagation through Monte Carlo (MC) estimates of the KL.

**Tractability of NFs** When optimizing normalizing flows with stochastic gradient descent, we have to evaluate a base density and compute the gradient with respect to its inputs. This poses no challenge as we have the flexibility to choose a simple density (e.g. uniform or Gaussian). In addition, for every  $x$  (i.e. an observation in parametric density estimation (Equation 2) or a base sample in variational inference (Equation 3)), the term  $|\det \mathbf{J}_{f_\theta(x)}|$  has to be evaluated and differentiated with respect to the parameters  $\theta$ . Note that  $f_\theta(x)$  is required to be invertible, as expressive as possible, and ideally fast to compute. In general, it is non-trivial to construct invertible transformations with efficiently computable  $|\det \mathbf{J}_{f_\theta(x)}|$ . Computing the determinant of a generic Jacobian  $\mathbf{J}_{f_\theta(x)} \in \mathbb{R}^{d \times d}$  runs in  $\mathcal{O}(d^3)$ -time. Our work and current research on NFs aim at constructing parametrized flows which meet efficiency requirements while maximizing the expressiveness of the densities they can represent.

### 2.2 AUTOREGRESSIVE FLOWS

We can construct  $f(x)$  such that its Jacobian is lower triangular, and thus has determinant  $\prod_{i=1}^d \partial f(x)_i / \partial x_i$ , which is computed in time  $\mathcal{O}(d)$ . Flows based on autoregressive transformations meet precisely this requirement (Kingma et al., 2016; Oliva et al., 2018; Huang et al., 2018). For a multivariate random variable  $X = \langle X_1, \dots, X_d \rangle$  with  $d > 1$ , we can use the chain rule to express the joint probability of  $x$  as product of  $d$  univariate conditional densities:

$$p_X(x) = p_{X_1}(x_1) \prod_{i=2}^d p_{X_i | X_{<i}}(x_i | x_{<i}). \quad (4)$$

When we then apply a normalizing flow to each univariate density, we have an autoregressive flow. Specifically, we can use a set of functions  $f^{(i)}$  that can be decomposed via *conditioners*  $c^{(i)}$ , and invertible *transformers*  $t^{(i)}$ :

$$y_i = f_\theta^{(i)}(x_{<i}) = t_\theta^{(i)}(x_i, c_\theta^{(i)}(x_{<i})), \quad (5)$$

where each transformer  $t^{(i)}$  must be an invertible function with respect to  $x_i$ , and each  $c^{(i)}$  is an unrestricted function. The resulting flow has a lower triangular Jacobian since each  $y_i$  depends only on  $x_{\leq i}$ . The flow is invertible when the Jacobian is constructed to have non-zero diagonal entries.

### 2.3 NEURAL AUTOREGRESSIVE FLOW

The invertibility of the flow as a whole depends on each  $t^{(i)}$  being an invertible function of  $x_i$ . For example, [Dinh et al. \(2014\)](#) and [Kingma et al. \(2016\)](#) model each  $t^{(i)}$  as an affine transformation whose parameters are predicted by  $c^{(i)}$ . As argued by [Huang et al. \(2018\)](#), these transformations were constructed to be trivially invertible, but their simplicity leads to a cap on expressiveness of  $f$ , thus requiring complex conditioners and a composition of multiple flows. They propose instead to learn a complex bijection using a neural network monotonic in  $x_i$  — this only requires constraining  $t^{(i)}$  to having non-negative weights and using strictly increasing activation functions. Figure 1a outlines a NAF. Each conditioner  $c^{(i)}$  is an unrestricted function of  $x_{<i}$ . To parametrize a monotonically increasing transformer network  $t^{(i)}$ , the outputs of each conditioner  $c^{(i)}$  are mapped to the positive real coordinate space by application of an appropriate activation (e.g. exp). The result is a flexible transformation with lower triangular Jacobian whose diagonal elements are positive.<sup>2</sup>


For efficient computation of all pseudo-parameters, as [Huang et al. \(2018\)](#) call the conditioners’ outputs, they use a masked autoregressive network ([Germain et al., 2015](#)). The Jacobian of a NAF is computed using the chain rule on  $f_\theta$  through all its hidden layers  $\{h^{(\ell)}\}_{\ell=1}^L$ :

$$\mathbf{J}_{f_\theta(x)} = \begin{bmatrix} \nabla_{h^{(L)}} y \end{bmatrix} \begin{bmatrix} \nabla_{h^{(L-1)}} h^{(L)} \end{bmatrix} \dots \begin{bmatrix} \nabla_x h^{(1)} \end{bmatrix}. \quad (6)$$


Since  $f_\theta$  is autoregressive,  $\mathbf{J}_{f_\theta(x)}$  is lower triangular and only the diagonal needs to be computed, i.e.  $\partial y_i / \partial x_i$  for each  $i$ . Thus, this operation requires only computing the derivatives of each  $t^{(i)}$ , reducing the time complexity.

Because the universal approximation theorem for densities holds for NAFs ([Huang et al., 2018](#)), increasing the expressiveness of a NAF is only a matter of employing larger transformer networks. However, the conditioner grows quadratically with the size of the transformer network and a combination of restricting the size of the transformer and a technique similar to *conditional weight normalization* ([Krueger et al., 2017](#)) is necessary to reduce the number of parameters. In §3 we propose to

<sup>2</sup>Note that the expressiveness of a NAF comes at the cost of analytic invertibility, i.e. though each  $t^{(i)}$  is bijective, thus invertible in principle, inverting the network itself is non-trivial.



(a) NAF: each  $c^{(i)}$  is a neural network that predicts pseudo-parameters for  $t^{(i)}$ , which in turn processes  $x_i$ .



(b) Our B-NAF: we do not use conditioner networks, instead we learn the flow network directly. Some weights are strictly positive (solid lines), others have no constraints (dashed lines).

Figure 1: Main differences between NAF ([Huang et al., 2018](#)) and our B-NAF. Both networks are autoregressive and invertible since  $y_i$  is processed with a function  $t^{(i)}$  which is monotonically increasing with respect to  $x_i$  and there is an arbitrary dependence on  $x_{<i}$ .

parametrize the transformer network directly, i.e. without a conditioner network, by exploiting the fact that the monotonicity constraint only requires  $\partial y_i / \partial x_i > 0$ , and therefore, careful design of a single feed-forward network can directly realize a transformation that is both autoregressive (with unconstrained manipulation of  $x_{<i}$ ) and strictly monotone (with positive  $\partial y_i / \partial x_i$ ).

## 3 BLOCK NEURAL AUTOREGRESSIVE FLOW

In the spirit of NAFs, we model each  $f_\theta^{(i)}(x_{\leq i})$  as a neural network with parameters  $\theta$ , but differently from NAFs, we do not predict  $\theta$  using a conditioner network, and instead, we learn  $\theta$  directly. In dense layers of  $f_\theta^{(i)}$ , we employ affine transformations with strictly positive weights to process  $x_i$ . This ensures strict monotonicity and thus invertibility of each  $f_\theta^{(i)}$  with respect to  $x_i$ . However, we do not impose this constraint on affine transformations of  $x_{<i}$ . Additionally, we need to always use invertible activation functions to ensure that the whole network is bijective (e.g., tanh or LeakyReLU). Each  $f_\theta^{(i)}$  is then a univariate flow implemented as an

arbitrarily wide and deep neural network which can approximate any invertible transformation. Much like other AFs, we can efficiently compute all  $f_\theta^{(i)}$  in parallel by employing a single masked autoregressive network (Germain et al., 2015). In the next section, we show how to construct each affine transformation using block matrices. From now on, we will refer to our novel family of flows as *block neural autoregressive flows* (B-NAFs).

### 3.1 AFFINE TRANSFORMATIONS WITH BLOCK MATRICES

For each affine transformation of  $x$ , we parametrize the bias term freely and we construct the weight matrix  $W \in \mathbb{R}^{ad \times bd}$  as a lower triangular *block* matrix for some  $a, b \geq 1$ . We use  $d \times (d+1)/2$  blocks  $B_{ij} \in \mathbb{R}^{a \times b}$  for  $i \in \{1, \dots, d\}$  and  $1 \leq j \leq i$ . We let  $B_{ij}$  (with  $i > j$ ) be freely parametrized and constrain diagonal blocks to be strictly positive applying an element-wise function  $g : \mathbb{R} \rightarrow \mathbb{R}_{>0}$  to each of them. Thus:

$$W = \begin{bmatrix} g(B_{11}) & 0 & \dots & 0 \\ B_{21} & g(B_{22}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ B_{d1} & B_{d2} & \dots & g(B_{dd}) \end{bmatrix}, \quad (7)$$

where we chose  $g(\cdot) = \exp(\cdot)$ . Since the flow has to preserve the input dimensionality, the first and the last affine transformations in the network must have  $b = 1$  and  $a = 1$ , respectively. Inside the network, the size of  $a$  and  $b$  can grow arbitrarily.

The intuition behind the construction of  $W$  is that every row of blocks  $B_{i1}, \dots, B_{ii}$  is a set of affine transformations (projections) that are processing  $x_{\leq i}$ . In particular, blocks in the upper triangular part of  $W$  are set to zero to make the flow autoregressive. Since the blocks  $B_{ii}$  are mapped to  $\mathbb{R}_{>0}$  through  $g$ , each transformation in such set is strictly monotonic for  $x_i$  and unconstrained on  $x_{< i}$ .

**B-NAF with masked networks** In practice, a more convenient parameterization of  $W$  consists of using a full matrix  $\hat{W} \in \mathbb{R}^{ad \times bd}$  which is then transformed applying two *masking* operations. One mask  $M_d \in \{0, 1\}^{ad \times bd}$  selects only elements in the diagonal blocks, and a second one  $M_o$  selects only off-diagonal and lower diagonal blocks. Thus, for each layer  $\ell$  we get

$$W^{(\ell)} = g(\hat{W}^{(\ell)}) \odot M_d^{(\ell)} + \hat{W}^{(\ell)} \odot M_o^{(\ell)}, \quad (8)$$

where  $\odot$  is the element-wise product. Figure 1b shows an outline of our block neural autoregressive flow.

Since each weight matrix  $W^{(\ell)}$  has some strictly positive and some zero entries, we need to take care of a

proper initialization which should take that into account. Indeed, weights are usually initialized to have a zero centred normally distributed output with variance dependent on the output dimensionality (Glorot and Bengio, 2010). Instead of carefully designing a new initialization technique to take care of this, we choose to initialize all blocks with a simple distribution and to apply weight normalization (Salimans and Kingma, 2016) to better cope with the effect of such initialization. See Appendix C for more details.

When constructing a stacked flow through a composition of  $n$  B-NAF transformations, we add gated residual connections for improving stability such that the composition is  $\hat{f}_n \circ \dots \circ \hat{f}_2 \circ \hat{f}_1$  where  $\hat{f}_i(x) = \alpha f_i(x) + (1 - \alpha)x$  and  $\alpha \in (0, 1)$  is a trainable scalar parameter.

### 3.2 AUTOREGRESSIVENESS AND INVERTIBILITY

In this section, we show that our flow  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$  meets the following requirements: i) its Jacobian  $\mathbf{J}_{f_\theta(x)}$  is lower triangular (needed for efficiency in computing its determinant), and ii) the diagonal entries of such Jacobian are positive (to ensure that  $f_\theta$  is a bijection).

**Proposition 1.** *The final Jacobian  $\mathbf{J}_{f_\theta(x)}$  of such transformation is lower triangular.*

*Proof sketch.* When applying the chain rule (Equation 6), the Jacobian of each affine transformation is  $W$  (Equation 8), a lower triangular block matrix, whereas the Jacobian of each element-wise activation function is a diagonal matrix. A matrix multiplication between a lower triangular block matrix and a diagonal matrix yields a lower triangular block matrix, and a multiplication between two lower triangular block matrices results in a lower triangular block matrix. Therefore, after multiplying all matrices in the chain, the overall Jacobian is lower triangular.  $\square$

**Proposition 2.** *When using strictly increasing activation functions (e.g., tanh or LeakyReLU), the diagonal entries of  $\mathbf{J}_{f_\theta(x)}$  are strictly positive.*

*Proof sketch.* When applying the chain rule (Equation 6), the Jacobian of each affine transformation has strictly positive values in its diagonal blocks where the Jacobian of each element-wise activation function is a diagonal matrix with strictly positive elements. When using matrix multiplication between two lower triangular block matrices (or one diagonal and one lower triangular block matrix)  $C = AB$  the resulting blocks on the diagonal of  $C$  are the result of a multiplication between only diagonal blocks of  $A, B$ . Indeed, such resulting blocks depend only on blocks of the same row and column partition. Using the notation of Equation 7, the resulting diagonal blocks of  $C$  are  $B_{ii}^{(C)} = g(B_{ii}^{(A)})g(B_{ii}^{(B)})$ . Therefore,



they are always positive. Eventually, using the chain rule, the final Jacobian is a lower triangular matrix with strictly positive elements in its diagonal.  $\square$

### 3.3 EFFICIENT JACOBIAN COMPUTATION

Proposition 1 is particularly useful for an efficient computation of  $\det \mathbf{J}_{f_\theta(x)}$  since we only need the product of its diagonal elements  $\partial y_i / \partial x_i$ . Thus, we can avoid computing the other entries. Since the determinant is the result of a product of positive values, we also remove the absolute-value operation resulting in

$$\log |\det \mathbf{J}_{f_\theta(x)}| = \sum_{i=0}^d \log (\mathbf{J}_{f_\theta(x)})_{ii}. \quad (9)$$

Additionally, as per Proposition 2, when using matrix multiplication, elements in the diagonal blocks (or entries) depend only on diagonal blocks of the same row and column partition. Since all diagonal blocks/entries are positive, we compute them directly in the log-domain to have more numerically stable operations:

$$\begin{aligned} \log (\mathbf{J}_{f_\theta(x)})_{ii} &= \log g(B_{ii}^{(\ell)}) \star \\ \log \mathbf{J}_{\sigma^{(\ell)}(h_\alpha^{(\ell-1)})} &\star \cdots \star \log g(B_{ii}^{(1)}), \end{aligned} \quad (10)$$

where  $\star$  denotes the log-matrix multiplication,  $\sigma^{(\ell)}$  the strictly increasing non-linear activation function at layer  $\ell$ , and  $\alpha$  indicates the set of indices corresponding to diagonal elements that depend on  $x_i$ . Notice that, since we chose  $g(\cdot) = \exp(\cdot)$  we can remove all redundant operations  $\log g(\cdot)$ . The log-matrix multiplication  $C = A \star B$  of two matrices  $A \in \mathbb{R}^{m \times n}$  and  $B \in \mathbb{R}^{n \times p}$  is  $C \in \mathbb{R}^{m \times p}$  where such operation can be implemented with a stable *log-sum-exp* operation (see details in Appendix D):

$$C_{ij} = \log \sum_{k=1}^n \exp(A_{ik} + B_{kj}). \quad (11)$$

A similar idea is also employed in NAF.

### 3.4 UNIVERSAL DENSITY APPROXIMATOR

In this section, we expose an intuitive proof sketch that our block neural autoregressive flow can approximate any real continuous probability density function (PDF).

Given a multivariate real and continuous random variable  $X = \langle X_1, \dots, X_d \rangle$ , its joint distribution can be factorized into a set of univariate conditional distributions (as in Equation 4), using an arbitrary ordering of the variables, and we can define a set of univariate conditional cumulative distribution functions (CDFs)  $Y_i =$

$F_{X_i|X_{<i}}(x_i|x_{<i}) = \mathbb{P}[X_i \leq x_i | X_{<i} = x_{<i}]$ . According to Hyvärinen and Pajunen (1999), such decomposition exists and each individual  $Y_i$  is independent as well as uniformly distributed in  $[0, 1]$ . Therefore, we can see  $F_X$  as a particular *normalizing flow* that maps  $X \in \mathbb{R}^n$  to  $Y \in [0, 1]^n$  where the distribution  $p_Y$  is uniform in the hyper-cube  $[0, 1]^n$ . Note that  $F_X$  is an autoregressive function and its Jacobian has a positive diagonal since  $\partial y_i / \partial x_i = p_{X_i|X_{<i}}(x_i|x_{<i})$ .

If each univariate flow  $f_\theta^{(i)}$  (see Equation 5) can approximate any invertible univariate conditional CDF, then  $f_\theta$  can approximate any PDF (Huang et al., 2018). Note that in general, a CDF  $F_{X_i|X_{<i}}$  is non-decreasing, thus not necessary invertible (Park and Park, 2018). Using B-NAF, each CDF is approximated with an arbitrarily large neural network and the output can be eventually mapped to  $(0, 1)$  with a sigmoidal function. Recalling that we only use positive weights for processing  $x_i$ , a neural network with non-negative weights is an universal approximator of monotonic functions (Daniels and Velikova, 2010). We use *strictly* positive weights to approximate a *strictly* monotonic function for  $x_i$  and we use arbitrary weights for  $x_{<i}$  (as there is no monotonicity constraint for them). Therefore, B-NAF can approximate any invertible CDF, and thus its corresponding PDF.

## 4 RELATED WORK

Current research on NFs focuses on constructing expressive parametrized invertible transformations with tractable Jacobians. Rezende and Mohamed (2015) were the first to employ parameterized flows in the context of variational inference proposing two parametric families: the planar and the radial flow. A drawback and bottleneck of such flows is that their power comes from stacking a large number of such transformations. More recently, van den Berg et al. (2018) generalized the use of planar flows showing improvements without increasing the number of transformations, and instead, by making each transformation more expressive.

In the context of density estimation, Germain et al. (2015) proposed MADE, a masked feed-forward network that efficiently computes an autoregressive transformation. MADEs are important building blocks in AFs, such as the inverse autoregressive flows (IAFs) introduced by Kingma et al. (2016). IAFs are based on trivially invertible affine transformations of the preceding coordinates of the input vector. The parameters of each transformation (a location and a positive scale) are predicted in parallel with a MADE, and therefore IAFs have a lower triangular Jacobian whose determinant is fast to evaluate. IAFs extend the parametric families




Figure 2: Comparison between Glow and B-NAF on density estimation for 2D toy data.

available for approximate posterior inference in variational inference. Neural autoregressive flow (NAF) by Huang et al. (2018) extend IAFs by generalizing the bijective transformation to one that can approximate any monotonically increasing function. They have been used both for parametric density estimation and approximate inference. In §2.3 we explain NAFs in detail and contrast them with our proposed B-NAFs (also, see Figure 1).

Larochelle and Murray (2011) were among the first to employ neural networks for autoregressive density estimation (NADE), in particular, for high-dimensional binary data. Non-linear independent components estimation (NICE) explored the direction of learning a map from high-dimensional data to a latent space with a simpler factorized distribution (Dinh et al., 2014). Papamakarios et al. (2017) proposed masked autoregressive flows (MAFs) as a generalization of real non-volume-preserving flows (Real NVP) by Dinh et al. (2017) showing improvements on density estimation.

In this work, we are modelling a discrete normalizing flow since at each transformation a discrete step is made. Continuous normalizing flows (CNF) were proposed by Chen et al. (2018) and modelled through a network that instead of predicting the output of the transformation predicts its derivative. The resulting transformation is computed using an ODE solver. Grathwohl et al. (2019) further improved such formulation proposing free-form Jacobian of reversible dynamics (FFJORD).

Orthogonal work has been done for constructing powerful invertible function such as invertible  $1 \times 1$  convolution (Glow) by Kingma and Dhariwal (2018), invertible  $d \times d$  convolutions (Hoogeboom et al., 2019), and invertible residual networks (Behrmann et al., 2019). Additionally,




Figure 3: Comparison between planar flow (PF) and B-NAF on four 2D energy functions from Table 1 of (Rezende and Mohamed, 2015).

Oliva et al. (2018) investigated different possibilities for the *conditioner* of an autoregressive transformation (e.g., recurrent neural networks).

## 5 EXPERIMENTS

### 5.1 DENSITY ESTIMATION ON TOY 2D DATA

In this experiment, we use our B-NAF to perform density estimation on 2-dimensional data as this helps us visualize the model capabilities to learn. We use the same toy data as Grathwohl et al. (2019) comparing the results with Glow (Kingma and Dhariwal, 2018), as they do. Given samples from a dataset with empirical distribution  $p_{\text{data}}$ , we parametrize a density  $p_{X|\theta}$  with a normalizing flow  $p_{X|\theta}(x) = p_Y(f_\theta(x)) |\det \mathbf{J}_{f_\theta(x)}|$  using B-NAF with  $p_Y$  a standard Normal distribution. We train for 20k iterations a single flow of B-NAF with 3 hidden layers of 100 units each using maximum likelihood estimation (i.e., maximizing  $\mathbb{E}_{p_{\text{data}}} [\log p_{X|\theta}(x)]$ , see Appendix A for more details and derivation of the objective). We used Adam optimizer (Kingma and Ba, 2014) with an initial learning rate of  $\alpha = 10^{-1}$  (and decay of 0.5 with patience of 2k steps), default  $\beta_1, \beta_2$ , and a batch size of 200. We took figures of Glow from (Grathwohl et al., 2019) who trained such models with 100 layers.

**Results** The learned distributions of both Glow and our method can be seen in Figure 2. Glow is capable of learn-

Table 1: Log-likelihood on the test set (higher is better) for 4 datasets (Dua and Karra Taniskidou, 2017) from UCI machine learning and BSDS300 (Martin et al., 2001). Here  $d$  is the dimensionality of datapoints and  $N$  the size of the dataset. We report average ( $\pm$ std) across 3 independently trained models. We also report the ratios between the number of parameters used by NAF-DDSF (with 5 or 10 flows) and our B-NAF. In highly dimensional datasets B-NAF uses orders of magnitude fewer parameters than NAF.

Model	POWER $\uparrow$	GAS $\uparrow$	HEPMASST $\uparrow$	MINIBOONE $\uparrow$	BSDS300 $\uparrow$
	$d=6;N=2,049,280$	$d=8;N=1,052,065$	$d=21;N=525,123$	$d=43;N=36,488$	$d=63;N=1,300,000$
Real NVP	0.17 $\pm$ .01	8.33 $\pm$ .14	-18.71 $\pm$ .02	-13.55 $\pm$ .49	153.28 $\pm$ 1.78
Glow	0.17 $\pm$ .01	8.15 $\pm$ .40	-18.92 $\pm$ .08	-11.35 $\pm$ .07	155.07 $\pm$ .03
MADE MoG	0.40 $\pm$ .01	8.47 $\pm$ .02	-15.15 $\pm$ .02	-12.27 $\pm$ .47	153.71 $\pm$ .28
MAF-affine	0.24 $\pm$ .01	10.08 $\pm$ .02	-17.73 $\pm$ .02	-12.24 $\pm$ .45	155.69 $\pm$ .28
MAF-affine MoG	0.30 $\pm$ .01	9.59 $\pm$ .02	-17.39 $\pm$ .02	-11.68 $\pm$ .44	156.36 $\pm$ .28
FFJORD	0.46 $\pm$ .01	8.59 $\pm$ .12	-14.92 $\pm$ .08	-10.43 $\pm$ .04	157.40 $\pm$ .19
NAF-DDSF	0.62 $\pm$ .01	11.96 $\pm$ .33	-15.09 $\pm$ .40	-8.86 $\pm$ .15	157.43 $\pm$ .30
TAN	0.60 $\pm$ .01	12.06 $\pm$ .02	-13.78 $\pm$ .02	-11.01 $\pm$ .48	159.80 $\pm$ .07
<b>Ours</b>	0.61 $\pm$ .01	12.06 $\pm$ .09	-14.71 $\pm$ .38	-8.95 $\pm$ .07	157.36 $\pm$ .03
Parameters ratio NAF (5) / B-NAF	2.29 $\times$	1.30 $\times$	17.94 $\times$	43.97 $\times$	8.24 $\times$
Parameters ratio NAF (10) / B-NAF	4.57 $\times$	2.60 $\times$	35.88 $\times$	87.91 $\times$	16.48 $\times$

ing a multi-modal distribution, but it has issues assigning the correct density in areas of low probability between disconnected regions. Our model is instead able to perfectly capture both multi-modality and discontinuities.

## 5.2 DENSITY MATCHING ON TOY 2D DATA

In this experiment, we use B-NAF to perform density matching on 2-dimensional target energy functions to visualize the model capabilities of matching them. We use the same energy functions described by Rezende and Mohamed (2015) comparing the results with them (using planar flows). For this task, we train a parameterized flow minimizing the KL divergence between the learned  $q_{Y|\theta}$  and the given target  $p_Y$ . We used a single flow using a B-NAF with 2 hidden layers of 100 units each. We train by minimizing  $\mathbb{KL}(q_{Y|\theta}||p_Y)$  (see Appendix B for a detailed derivation) using Monte Carlo sampling. We optimized using Adam for 20k iterations with an initial learning rate of  $\alpha = 10^{-2}$  (and decay of 0.5 with patience of 2k steps), default  $\beta_1, \beta_2$ , and a batch size of 200. Planar flow figures were taken from Chen et al. (2018). Note that planar flows were trained for 500k iterations using RMSProp (Tieleman and Hinton, 2012).

**Results** Figure 3 shows that our model perfectly matches all target distributions. Indeed, on functions 3 and 4 it looks like B-NAF can better learn the density in certain areas. The model capacity of planar normalizing flows is determined by their depth ( $K$ ) and Rezende and Mohamed (2015) had to stack 32 flows to match the

energy function reasonably well. Deeper networks are harder to optimize, and our flow matches all the targets using a neural network with only 2 hidden layers.

## 5.3 REAL DATA DENSITY ESTIMATION

In this experiment, we use a B-NAF to perform density estimation on 5 real datasets. Similarly to Section 5.1, we train using MLE maximizing  $\mathbb{E}_{p_{\text{data}}}[\log p_{X|\theta}(x)]$ . We compare our results against Real NVP (Dinh et al., 2017), Glow (Kingma and Dhariwal, 2018), MADE (Germain et al., 2015), MAF (Papamakarios et al., 2017), TAN (Oliva et al., 2018), NAF (Huang et al., 2018), and FFJORD (Grathwohl et al., 2019). For our B-NAF, we stacked 5 flows and we employed a small grid search on the number of layers and the size of hidden units per flow ( $L \in \{1, 2\}$  and  $H \in \{10d, 20d, 40d\}$ , respectively, where  $d$  is the input size of datapoints which is different for each dataset). When stacking B-NAF flows, the elements of each output vector are permuted so that a different set of elements is considered at each flow. This technique is not novel and it is also used by Dinh et al. (2017); Papamakarios et al. (2017); Kingma et al. (2016). We trained using Adam with Polyak averaging (with  $\phi = 0.998$ ) as in NAF (Polyak and Juditsky, 1992). We also applied an exponentially decaying learning rate schedule (from  $\alpha = 10^{-2}$  with rate  $\lambda = 0.5$ ) based on no-improvement with patience of 20 epochs. We trained until convergence (but maximum 1k epochs), stopping after 100 epochs without improvement on validation set.

**Datasets** Following Papamakarios et al. (2017), we perform unconditional density estimation on four datasets (Dua and Karra Taniskidou, 2017) from UCI machine learning repository<sup>3</sup> as well as one dataset of patches of images (Martin et al., 2001).

**Results** Table 1 shows the results of this experiment reporting log-likelihood on test set. In all datasets, our B-NAF is better than Real NVP, Glow, MADE, and MAF and it performs comparable or better to NAF. B-NAF also outperforms FFJORD in all dataset except on BSDS300 where there is a marginal difference ( $< 0.02\%$ ) between the two methods. On GAS and HEPMASS, B-NAF performs better than most of the other models and even better than NAF. In the other datasets, the gap in performance compared to NAF is marginal. We observed that in most datasets, the best performing model was the largest one in the grid search. It is possible that we do a too narrow hyper-parameter search compared to what other methods do. For instance, FFJORD results come from a wider grid search than ours where Grathwohl et al. (2019), Huang et al. (2018), and Oliva et al. (2018) varied the number of flows during tuning.

We compare NAF and our B-NAF in terms of the number of parameters employed and report the ratio between the two for each dataset in Table 1 (bottom). For datasets with low-dimensional datapoints (i.e, GAS and POWER) our model uses a comparable number of parameters to NAF. For high-dimensional datapoints the gap between the parameters used by NAF and B-NAF grows, with B-NAF much smaller, as we intended. For instance, on both HEPMASS and MINIBOONE, our models have marginal differences in performance with NAF while having respectively  $\sim 18\times$  and  $\sim 40\times$  fewer parameters than NAF. This evidence supports our argument that NAF models are over-parametrized and it is possible to achieve similar performance with an order of magnitude fewer parameters. Besides, when training models on GPUs, being memory efficient allows to train more models in parallel on the same device. Additionally, in general, a normalizing flow can be a component of a larger architecture that might require more memory than the flow itself (as the models for experiments in the next Section).

## 5.4 VARIATIONAL AUTO-ENCODERS

An interesting application of our framework is modelling more flexible posterior distributions in a variational auto-encoder (VAE) setting (Kingma and Welling, 2013). In this setting, we assume that an observation  $x$  (i.e., the data) is drawn from the marginal of a deep la-

tent variable model, i.e.  $X \sim p_{X|\theta}$ , where  $p_{X|\theta}(x) = \int p_Z(z)p_{X|Z,\theta}(x|z)dz$  where  $Z \sim \mathcal{N}(0, I)$  is unobserved. The goal is performing maximum likelihood estimation of the marginal. Since  $Z$  is not observed, maximizing the objective would require marginalization over the latent variables, which is generally intractable. Using variational inference (Jordan et al., 1999), we can maximize a lower bound on log-likelihood:

$$\log p_{X|\theta}(x) \geq \mathbb{E}_{q_{Z|X,\phi}(z)} \left[ \log \frac{p_{XZ|\theta}(x, z)}{q_{Z|X,\phi}(z|x)} \right], \quad (12)$$

where  $p_{X|Z,\theta}$  and  $q_{Z|X,\phi}$  are parametrized via neural networks with learnable parameters  $\theta$  and  $\phi$  (Kingma and Welling, 2013), in particular,  $q_{Z|X,\phi}$  is an approximation to the intractable posterior  $p_{Z|X,\theta}$ . This bound is called the evidence lower bound (ELBO), and maximizing the ELBO is equivalent to minimizing  $\mathbb{KL}(q_{Z|X,\phi}||p_{Z|X,\theta})$ . The more expressive the approximating family is, the more likely we are to obtain a tight bound. Recent literature approaches tighter bounds by approximating the posterior with normalizing flows. Also note that NFs reparametrize  $q_{Z|X,\phi}(z|x) = q_Y(f_\phi(z; x))|\det \mathbf{J}_{f_\phi(z;x)}|$  via a simpler fixed base distribution, e.g. a standard Gaussian, and therefore we can follow stochastic gradient estimates of the ELBO with respect to both sets of parameters. In this experiment, we use our flow for posterior approximation showing that B-NAF compares with recently proposed NFs for variational inference. We reproduce experiments by van den Berg et al. (2018) (Sylvester flows or SNF) while replacing their flow with ours. We keep the encoder and decoder networks exactly the same to fairly compare with all models trained with such procedure. We compare our B-NAF to their flows on the same 4 datasets as well as to a normal VAE (Kingma and Welling, 2013), planar flows (Rezende and Mohamed, 2015), and IAFs (Kingma et al., 2016).<sup>4</sup>

In this experiment, the input dimensionality of the flow is fixed to  $d = 64$ . We employed a small grid search on the MNIST dataset for the number of flows  $K \in \{4, 8\}$ , and for the number of hidden units per flow  $H \in \{2d, 4d, 8d\}$  while keeping the number of layers fixed at  $L = 1$ . The elements of each output vector are permuted after each B-NAF flow (as we do in § 5.3). We keep the best hyper-parameters of this search for the other datasets. We train using Adamax with  $\alpha = 5 \cdot 10^{-4}$ . We refer to Appendix A of van den Berg et al. (2018) for details on the network architectures for the encoder and decoder.

**Datasets** Following van den Berg et al. (2018) we carried our experiments on 4 datasets: statically binarized

<sup>3</sup><http://archive.ics.uci.edu/ml>

<sup>4</sup> Although Huang et al. (2018) also experimented with VAEs using NAF, they used only one dataset (MNIST) and employed a different encoder/decoder architecture than van den Berg et al. (2018). Therefore, results are not comparable.



Table 2: Negative log-likelihood (NLL) and negative evidence lower bound (-ELBO) for static MNIST, Freyfaces, Omniglot and Caltech 101 Silhouettes datasets. For the Freyfaces dataset the results are reported in bits per dim. For the other datasets the results are reported in nats. For all datasets we report the mean and the standard deviations over 3 runs with different random initializations.

Model	MNIST		Freyfaces		Omniglot		Caltech 101	
	-ELBO↓	NLL↓	-ELBO↓	NLL↓	-ELBO↓	NLL↓	-ELBO↓	NLL↓
VAE	86.55±.06	82.14±.07	4.53±.02	4.40±.03	104.28±.39	97.25±.23	110.80±.46	99.62±.74
Planar	86.06±.31	81.91±.22	4.40±.06	4.31±.06	102.65±.42	96.04±.28	109.66±.42	98.53±.68
IAF	84.20±.17	80.79±.12	4.47±.05	4.38±.04	102.41±.04	96.08±.16	111.58±.38	99.92±.30
Sylvester	83.32±.06	80.22±.03	4.45±.04	4.35±.04	99.00±.04	93.77±.03	104.62±.29	93.82±.62
<b>Ours</b>	83.59±.15	80.71±.09	4.42±.05	4.33±.04	100.08±.07	94.83±.10	105.42±.49	94.91±.51

MNIST (Larochelle and Murray, 2011), Freyfaces,<sup>5</sup> Omniglot (Lake et al., 2015) and Caltech 101 Silhouettes (Marlin et al., 2010). All those datasets consist of black and white images of different sizes.

**Amortizing flow parameters** When using NFs in an amortized inference setting, the parameters of each flow are not learned directly but predicted with another function from each datapoint (Rezende and Mohamed, 2015). In our case, we do not amortize all parameters of B-NAF since that would require very large predictors and we want to keep our flow memory efficient. Alternatively, every affine matrix  $W \in \mathbb{R}^{n \times m}$  is shared among all datapoints. Then, for each affine transformation, we achieve a degree of amortization by predicting 3 vectors, the bias  $b \in \mathbb{R}^n$  and 2 vectors  $v_1 \in \mathbb{R}^n$  and  $v_2 \in \mathbb{R}^m$  that we multiply row- and column-wise respectively to  $W$ .

**Results** Table 2 shows the results of these experiments. From the grid search, it turned out that the best B-NAF model has  $K = 8$  (flows) and  $H = 4d$  (hidden units). Note that the best models reported by van den Berg et al. (2018) used 16 flows. Our model is quite flexible without being as deep as other models. Results show that B-NAF is better than normal VAE, planar flows, and IAFs in all four datasets. Although B-NAF performs slightly worse than Sylvester flows, van den Berg et al. (2018) applied a full amortization for the parameters of the flow, while we do not. They proposed two alternative parametrizations to construct Sylvester flows: orthogonal SNF and Householder SNF. For each datapoint, SNF has to predict from 50.7k to 76.8k values (depending on the parametrization) to fully amortize parameters of the flow, while we use only 7.7k (i.e., 6.64× to 10.0× fewer). Notably, recalling that these are not trainable parameters, we use 6.16× (orthogonal SNF) and 9.35× (Householder SNF) fewer trainable parameters as well. Besides, we also use

<sup>5</sup>[http://www.cs.nyu.edu/~roweis/data/frey\\_rawface.mat](http://www.cs.nyu.edu/~roweis/data/frey_rawface.mat)

14.45× fewer parameters than IAF. This shows that IAF and SNF are over-parametrized too, and it is possible to achieve similar performance in the context of variational inference with an order of magnitude fewer parameters.

## 6 CONCLUSIONS

We present a new family of flexible normalizing flows, *block neural autoregressive flows*. B-NAFs are universal approximators of density functions and maintain an efficient parametrization. Our flow is based on directly parametrizing a transformation that guarantees autoregressiveness and monotonicity without the need for large conditioner networks and without compromising parallelism. Compared to the only other flow (to the best of our knowledge) which is also a universal approximator, our B-NAFs require orders of magnitudes fewer parameters. We validate B-NAFs on parametric density estimation on toy and real datasets, as well as, on approximate posterior inference for deep latent variable models, showing favorable performance across datasets and against various established flows. For future work, we are interested in at least two directions. One concerns gaining access to the inverse of the flow—note that, while B-NAFs and NAFs are invertible in principle, their inverses are not available in closed form. Another concerns deep generative models with large decoders (e.g. in natural language processing applications): since we achieve high flexibility at a low memory footprint our flows seem to be a good fit.

## Acknowledgements

This project is supported by SAP Innovation Center Network, ERC Starting Grant BroadSem (678254) and the Dutch Organization for Scientific Research (NWO) VIDI 639.022.518. Wilker Aziz is supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 825299 (Gourmet).

## References

- Behrmann, J., Duvenaud, D., and Jacobsen, J.-H. (2019). Invertible residual networks. *Proceedings of the International Conference on Machine Learning (ICML)*.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. (2018). Neural ordinary differential equations. *Advances in Neural Information Processing Systems*.
- Daniels, H. and Velikova, M. (2010). Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6):906–917.
- Dinh, L., Krueger, D., and Bengio, Y. (2014). Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2017). Density estimation using real NVP. *Proceedings of the 5th International Conference on Learning Representations (ICLR)*.
- Dua, D. and Karra Taniskidou, E. (2017). UCI machine learning repository.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. (2015). Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pages 881–889.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. (2019). FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. *International Conference on Learning Representations*.
- Ha, D., Dai, A., and Le, Q. V. (2017). Hypernetworks. *International Conference on Learning Representations*.
- Hoogeboom, E., Berg, R. v. d., and Welling, M. (2019). Emerging convolutions for generative normalizing flows. *Proceedings of the International Conference on Machine Learning (ICML)*.
- Huang, C.-W., Krueger, D., Lacoste, A., and Courville, A. (2018). Neural autoregressive flows. *International Conference on Learning Representations*.
- Hyvärinen, A. and Pajunen, P. (1999). Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Kingma, D. P. and Dhariwal, P. (2018). Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10236–10245.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *International Conference on Learning Representations*.
- Krueger, D., Huang, C.-W., Islam, R., Turner, R., Lacoste, A., and Courville, A. (2017). Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338.
- Larochelle, H. and Murray, I. (2011). The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 29–37.
- Marlin, B., Swersky, K., Chen, B., and Freitas, N. (2010). Inductive principles for restricted boltzmann machine learning. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 509–516.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. (2001). A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference*, volume 2, pages 416–423. IEEE.
- Oliva, J., Dubey, A., Zaheer, M., Poczos, B., Salakhutdinov, R., Xing, E., and Schneider, J. (2018). Transformation autoregressive networks. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3898–3907, Stockholmsmssan, Stockholm Sweden. PMLR.
- Papamakarios, G., Pavlakou, T., and Murray, I. (2017). Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347.

- Park, K. I. and Park (2018). *Fundamentals of Probability and Stochastic Processes with Applications to Communications*. Springer.
- Polyak, B. T. and Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855.
- Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 1530–1538. JMLR. org.
- Rippel, O. and Adams, R. P. (2013). High-dimensional probability estimation with deep density models. *arXiv preprint arXiv:1302.5125*.
- Salimans, T. and Kingma, D. P. (2016). Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 901–909.
- Tabak, E. G., Vanden-Eijnden, E., et al. (2010). Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233.
- Tieleman, T. and Hinton, G. (2012). Lecture 6.5-rmsprop, coursera: Neural networks for machine learning. *University of Toronto, Technical Report*.
- van den Berg, R., Hasenclever, L., Tomczak, J. M., and Welling, M. (2018). Sylvester normalizing flows for variational inference. *34th Conference on Uncertainty in Artificial Intelligence (UAI18)*.