
Unconstrained Monotonic Neural Networks

Antoine Wehenkel
University of Liège

Gilles Louppe
University of Liège

Abstract

Monotonic neural networks have recently been proposed as a way to define invertible transformations. These transformations can be combined into powerful autoregressive flows that have been shown to be universal approximators of continuous probability distributions. Architectures that ensure monotonicity typically enforce constraints on weights and activation functions, which enables invertibility but leads to a cap on the expressiveness of the resulting transformations. In this work, we propose the Unconstrained Monotonic Neural Network (UMNN) architecture based on the insight that a function is monotonic as long as its derivative is strictly positive. In particular, this latter condition can be enforced with a free-form neural network whose only constraint is the positiveness of its output. We evaluate our new invertible building block within a new autoregressive flow (UMNN-MAF) and demonstrate its effectiveness on density estimation experiments. We also illustrate the ability of UMNNs to improve variational inference.

1 Introduction

Monotonic neural networks have been known as powerful tools to build monotone models of a response variable with respect to individual explanatory variables [Archer and Wang, 1993, Sill, 1998, Daniels and Velikova, 2010, Gupta et al., 2016, You et al., 2017]. Recently, strictly monotonic neural networks have also been proposed as a way to define invertible transformations. These transformations can be combined into effective autoregressive flows that can be shown to be universal approximators of continuous probability distributions. Examples include Neural Autoregressive Flows [NAF, Huang et al., 2018] and Block Neural Autoregressive Flows [B-NAF, De Cao et al., 2019]. Architectures that ensure monotonicity typically enforce constraints on weight and activation functions, which enables invertibility but leads to a cap on the expressiveness of the resulting transformations. For neural autoregressive flows, this does not impede universal approximation but typically requires either complex conditioners or a composition of multiple flows.

Nevertheless, autoregressive flows defined as stacks of reversible transformations have proven to be quite efficient for density estimation of empirical distributions [Papamakarios et al., 2019, 2017, Huang et al., 2018], as well as to improve posterior modeling in Variational Auto-Encoders (VAE) [Germain et al., 2015, Kingma et al., 2016, Huang et al., 2018]. Practical successes of these models include speech synthesis [van den Oord et al., 2016, Oord et al., 2018], likelihood-free inference [Papamakarios et al., 2019], probabilistic programming [Tran et al., 2017] and image generation [Kingma and Dhariwal, 2018]. While stacking multiple reversible transformations improves the capacity of the full transformation to represent complex probability distributions, it remains unclear which class of reversible transformations should be used.

In this work, we propose a class of reversible transformations based on a new Unconstrained Monotonic Neural Network (UMNN) architecture. We base our contribution on the insight that a function is monotonic as long as its derivative is strictly positive. This latter condition can be enforced with a free-form neural network whose only constraint is for its output to remain strictly positive.

We summarize our contributions as follows:

- We introduce the Unconstrained Monotonic Neural Network (UMNN) architecture, a new reversible scalar transformation defined via a free-form neural network.
- We combine UMNN transformations into an autoregressive flow (UMNN-MAF) and we demonstrate competitive or state-of-the-art results on benchmarks for normalizing flows.
- We empirically illustrate the scalability of our approach by applying UMNN on high dimensional density estimation problems.

2 Unconstrained monotonic neural networks

Our primary contribution consists in a neural network architecture that enables learning arbitrary monotonic functions. More specifically, we want to learn a strictly monotonic scalar function $F(x; \psi) : \mathbb{R} \rightarrow \mathbb{R}$ without imposing strong constraints on the expressiveness of the hypothesis class. In UMNNs, we achieve this by only imposing the derivative $f(x; \psi) = \frac{\partial F(x; \psi)}{\partial x}$ to remain of constant sign or, without loss of generality, to be strictly positive. As a result, we can parameterize the bijective mapping $F(x; \psi)$ via its strictly positive derivative $f(x; \psi)$ as

$$F(x; \psi) = \int_0^x f(t; \psi) dt + \underbrace{F(0; \psi)}_{\beta}, \quad (1)$$

where $f(t; \psi) : \mathbb{R} \rightarrow \mathbb{R}_+$ is a strictly positive parametric function and $\beta \in \mathbb{R}$ is a scalar. We make f arbitrarily complex using an unconstrained neural network whose output is forced to be strictly positive through an ELU activation unit increased by 1. ψ denotes the parameters of this neural network.

Forward integration The forward evaluation of $F(x; \psi)$ requires solving the integral in Equation (1). While this might appear daunting, such integrals can often be efficiently approximated numerically using Clenshaw-Curtis quadrature. The better known trapezoidal rule, which corresponds to the two-point Newton-Cotes quadrature rule, has an exponential convergence when the integrand is periodic and the range of integration corresponds to its period. Clenshaw-Curtis quadrature takes advantage of this property by using a change of variables followed by a cosine transform. This extends the exponential convergence of the trapezoidal rule for periodic functions to any Lipschitz continuous function. As a result, the number of evaluation points required to reach convergence grows with the Lipschitz constant of the function.

Backward integration Training the integrand neural network f requires evaluating the gradient of F with respect to its parameters. While this gradient could be obtained by backpropagating directly through the integral solver, this would also result in a memory footprint that grows linearly with the number of integration steps. Instead, the derivative of an integral with respect to a parameter ω can be expressed with the Leibniz integral rule:

$$\frac{d}{d\omega} \left(\int_{a(\omega)}^{b(\omega)} f(t; \omega) dt \right) = f(b(\omega); \omega) \frac{d}{d\omega} b(\omega) - f(a(\omega); \omega) \frac{d}{d\omega} a(\omega) + \int_{a(\omega)}^{b(\omega)} \frac{\partial}{\partial \omega} f(t; \omega) dt. \quad (2)$$

Applying Equation (2) to evaluate the derivative of Equation (1) with respect to the parameters ψ , we find

$$\begin{aligned} \nabla_{\psi} F(x; \psi) &= f(x; \psi) \nabla_{\psi} (x) - f(0; \psi) \nabla_{\psi} (0) + \int_0^x \nabla_{\psi} f(t; \psi) dt + \nabla_{\psi} \beta \\ &= \int_0^x \nabla_{\psi} f(t; \psi) dt + \nabla_{\psi} \beta. \end{aligned} \quad (3)$$

When using a UMNN block in a neural architecture, it is also important to be able to compute its derivative with respect to its input x . In this case, applying Equation (2) leads to

$$\frac{d}{dx} F(x; \psi) = f(x; \psi). \quad (4)$$

Equations (3) and (4) make the memory footprint for the backward pass independent from the number of integration steps, and therefore also from the desired accuracy. Indeed, instead of computing the gradient of the integral (which requires keeping track of all the integration steps), we integrate the gradient (which is memory efficient, as this corresponds to summing gradients at different evaluation points). We provide the pseudo-code of the forward and backward passes using Clenshaw-Curtis quadrature in Appendix B.

Numerical inversion In UMMNs, the modeled monotonic function F is arbitrary. As a result, computing its inverse cannot be done analytically. However, since F is strictly monotonic, it admits a unique inverse x for any point $y = F(x; \psi)$ in its image, therefore inversion can be computed efficiently with common root-finding algorithms. In our experiments, search algorithms such as the bisection method proved to be fast enough.

3 UMNN autoregressive models

3.1 Normalizing flows

A Normalizing Flow [NF, [Rezende and Mohamed, 2015](#)] is defined as a sequence of invertible transformations $\mathbf{u}_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$ ($i = 1, \dots, k$) composed together to create an expressive invertible mapping $\mathbf{u} = \mathbf{u}_1 \circ \dots \circ \mathbf{u}_k : \mathbb{R}^d \rightarrow \mathbb{R}^d$. It is common for normalizing flows to stack the same parametric function \mathbf{u}_i (with different parameters values) and to reverse variables ordering after each transformation. For this reason, we will focus on how to build one of these repeated transformations, which we further refer to as $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

Density estimation NFs are most commonly used for density estimation, that map empirical samples to unstructured noise. Using normalizing flows, we define a bijective mapping $\mathbf{u}(\cdot; \theta) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ from a sample $\mathbf{x} \in \mathbb{R}^d$ to a latent vector $\mathbf{z} \in \mathbb{R}^d$ equipped with a density $p_Z(\mathbf{z})$. The transformation \mathbf{u} implicitly defines a density $p(\mathbf{x}; \theta)$ as given by the change of variables formula,

$$p(\mathbf{x}; \theta) = p_Z(\mathbf{u}(\mathbf{x}; \theta)) \left| \det J_{\mathbf{u}(\mathbf{x}; \theta)} \right|, \quad (5)$$

where $J_{\mathbf{u}(\mathbf{x}; \theta)}$ is the Jacobian of $\mathbf{u}(\mathbf{x}; \theta)$ with respect to \mathbf{x} . The resulting model is trained by maximizing the likelihood of the data $\{\mathbf{x}^1, \dots, \mathbf{x}^N\}$.

Variational auto-encoders NFs are also used in VAE to improve posterior modeling. In this case, a normalizing flow transforms a distribution p_Z into a complex distribution q which can better model the variational posterior. The change of variables formula yields

$$q(\mathbf{u}(\mathbf{z}; \theta)) = p_Z(\mathbf{z}) \left| \det J_{\mathbf{u}(\mathbf{z}; \theta)} \right|^{-1}. \quad (6)$$

3.2 Autoregressive transformations

To be of practical use, NFs must be composed of transformations for which the determinant of the Jacobian can be computed efficiently, otherwise its evaluation would be running in $\mathcal{O}(d^3)$. A common solution consists in making the transformation \mathbf{g} autoregressive, i.e., such that $\mathbf{g}(\mathbf{x}; \theta)$ can be rewritten as a vector of d scalar functions,

$$\mathbf{g}(\mathbf{x}; \theta) = [g^1(x_1; \theta) \quad \dots \quad g^i(\mathbf{x}_{1:i}; \theta) \quad \dots \quad g^d(\mathbf{x}_{1:d}; \theta)],$$

where $\mathbf{x}_{1:i} = [x_1 \quad \dots \quad x_i]^T$ is the vector including the i first elements of the full vector \mathbf{x} . The Jacobian of this function is lower triangular, which makes the computation of its determinant $\mathcal{O}(d)$. Enforcing the bijectivity of each component g^i is then sufficient to make \mathbf{g} bijective as well.

For the multivariate density $p(\mathbf{x}; \theta)$ induced by $\mathbf{g}(\mathbf{x}; \theta)$ and $p_Z(\mathbf{z})$, we can use the chain rule to express the joint probability of \mathbf{x} as a product of d univariate conditional densities,

$$p(\mathbf{x}; \theta) = p(x_1; \theta) \prod_{i=1}^{d-1} p(x_{i+1} | \mathbf{x}_{1:i}; \theta). \quad (7)$$

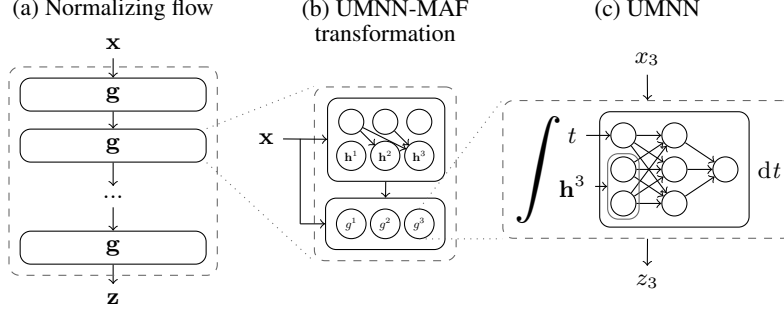


Figure 1: **(a)** A normalizing flow made of repeated UMNN-MAF transformations g with identical architectures. **(b)** A UMNN-MAF which transforms a vector $\mathbf{x} \in \mathbb{R}^3$. **(c)** The UMNN network used to map x_3 to z_3 conditioned on the embedding $\mathbf{h}^3(\mathbf{x}_{1:2})$.

When $p_Z(\mathbf{z})$ is a factored distribution $p_Z(\mathbf{z}) = \prod_{i=1}^d p(z_i)$, we identify that each component z_i coupled with the corresponding function g^i encodes for the conditional $p(x_i | \mathbf{x}_{1:i-1}; \theta)$. Autoregressive transformations strongly rely on the expressiveness of the scalar functions g^i . In this work, we propose to use UMNNs to create powerful bijective scalar transformations.

3.3 UMNN autoregressive transformations (UMNN-MAF)

We now combine UMNNs with an embedding of the conditioning variables to build invertible autoregressive functions g^i . Specifically, we define

$$\begin{aligned} g^i(\mathbf{x}_{1:i}; \theta) &= F^i(x_i, \mathbf{h}^i(\mathbf{x}_{1:i-1}; \phi^i); \psi^i) \\ &= \int_0^{x_i} f^i(t, \mathbf{h}^i(\mathbf{x}_{1:i-1}; \phi^i); \psi^i) dt + \beta^i(\mathbf{h}^i(\mathbf{x}_{1:i-1}; \phi^i)), \end{aligned} \quad (8)$$

where $\mathbf{h}^i(\cdot; \phi^i) : \mathbb{R}^{i-1} \rightarrow \mathbb{R}^q$ is a q -dimensional neural embedding of the conditioning variables $\mathbf{x}_{1:i-1}$ and $\beta(\cdot)^i : \mathbb{R}^{i-1} \rightarrow \mathbb{R}$. Both degenerate into constants for $g^1(\mathbf{x}_1)$. The parameters θ of the whole transformation $g(\cdot; \theta)$ is the union of all parameters ϕ^i and ψ^i . For simplicity we remove the parameters of the networks by rewriting $f^i(\cdot; \psi^i)$ as $f^i(\cdot)$ and $\mathbf{h}^i(\cdot; \phi^i)$ as $\mathbf{h}^i(\cdot)$.

In our implementation, we use a Masked Autoregressive Network [Germain et al., 2015, Kingma et al., 2016, Papamakarios et al., 2017] to simultaneously parameterize the d embeddings. In what follows we refer to the resulting UMNN autoregressive transformation as UMNN-MAF. Figure 1 summarizes the complete architecture.

Log-density The change of variables formula applied to the UMMN autoregressive transformation results in the log-density

$$\begin{aligned} \log p(\mathbf{x}; \theta) &= \log p_Z(g(\mathbf{x}; \theta)) |\det J_{g(\mathbf{x}; \theta)}| \\ &= \log p_Z(g(\mathbf{x}; \theta)) + \log \left| \prod_{i=1}^d \frac{\partial F^i(x_i, \mathbf{h}^i(\mathbf{x}_{1:i-1}))}{\partial x_i} \right| \\ &= \log p_Z(g(\mathbf{x}; \theta)) + \sum_{i=1}^d \log f^i(x_i, \mathbf{h}^i(\mathbf{x}_{1:i-1})). \end{aligned} \quad (9)$$

Therefore, the transformation leads to a simple expression of (the determinant of) its Jacobian, which can be computed efficiently with a single forward pass. This is different from FFJORD [Grathwohl et al., 2018] which relies on numerical methods to compute both the Jacobian and the transformation between the data and the latent space. Therefore our proposed method makes the computation of the Jacobian exact and efficient at the same time.

Sampling Generating samples require evaluating the inverse transformation $g^{-1}(\mathbf{z}; \theta)$. The components of the inverse vector $\mathbf{x}^{\text{inv}} = g^{-1}(\mathbf{z}; \theta)$ can be computed recursively by inverting each

component of $g(\mathbf{x}; \boldsymbol{\theta})$:

$$x_1^{\text{inv}} = (g^1)^{-1}(z_1; \mathbf{h}^1) \quad \text{if } i = 1 \quad (10)$$

$$x_i^{\text{inv}} = (g^i)^{-1}(z_i; \mathbf{h}^i(\mathbf{x}_{1:i-1}^{\text{inv}})) \quad \text{if } i > 1 \quad (11)$$

where $(g^i)^{-1}$ is the inverse of g^i . Another approach to invert an autoregressive model would be to approximate its inverse with another autoregressive network [Oord et al., 2018]. In this case, the evaluation of the approximated inverse model is as fast as the forward model.

Universality Since the proof is straightforward, we only sketch that UMNN-MAF is a universal density approximator of continuous random variables. We rely on the inverse sampling theorem to prove that UMNNs are universal approximators of continuously derivable (\mathbb{C}^1) monotonic functions. Indeed, if UMNNs can represent any \mathbb{C}^1 monotonic function, then they can also represent the (inverse) cumulative distribution function of any continuous random variable. Any continuously derivable function $f : \mathcal{D} \rightarrow \mathcal{I}$ can be expressed as the following integral: $f(x) = \int_a^x \frac{df}{dx} dx + f(a)$, $\forall x, a \in \mathcal{D}$. The derivative $\frac{df}{dx}$ is a continuous positive function and the universal approximation theorem of NNs ensures it can be successfully approximated with a NN of sufficient capacity (such as those used in UMNNs).

4 Related work

The most similar work to UMNN-MAF are certainly Neural Autoregressive Flow [NAF, Huang et al., 2018] and Block Neural Autoregressive Flow [B-NAF, De Cao et al., 2019], which both rely on strictly monotonic transformations for building bijective mappings. In NAF, transformations are defined as neural networks which activation functions are all constrained to be strictly monotonic and which weights are the output of a strictly positive and autoregressive HyperNetwork [Ha et al., 2017]. Huang et al. [2018] shows that NAFs are universal density approximators. In B-NAF, the authors improve on the scalability of the NAF architecture by making use of masking operations instead of HyperNetworks. They also present a proof of the universality of B-NAF, which extends to UMNN-MAF. Our work differs from both NAF and B-NAF in the sense that the UMNN monotonic transformation is based on free-form neural networks for which no constraint, beyond positiveness of the output, is enforced on the hypothesis class. This leads to multiple advantages: it enables the use of any state-of-the-art neural architecture, simplifies weight initialization, and leads to a more lightweight evaluation of the Jacobian.

More generally, UMNN-MAF relates to works on normalizing flows built upon autoregressive networks and affine transformations. Germain et al. [2015] first introduced masking as an efficient way to build autoregressive networks, and proposed autoregressive networks for density estimation of high dimensional binary data. Masked Autoregressive Flows [Papamakarios et al., 2017] and Inverse Autoregressive Flows [Kingma et al., 2016] have generalized this approach to real data, respectively for density estimation and for latent posterior representation in variational auto-encoders. More recently, Oliva et al. [2018] proposed to stack various autoregressive architectures to create powerful reversible transformations. Meanwhile, Jaini et al. [2019] proposed a new Sum-of-Squares flow that is defined as the integral of a second order polynomial parametrized by an autoregressive NN.

With NICE, Dinh et al. [2015] introduced coupling layers, which correspond to bijective transformations splitting the input vector into two parts. They are defined as

$$\mathbf{z}_{1:k} = \mathbf{x}_{1:k} \quad \text{and} \quad \mathbf{z}_{k+1:d} = e^{\boldsymbol{\sigma}(\mathbf{x}_{1:k})} \odot \mathbf{x}_{k+1:d} + \boldsymbol{\mu}(\mathbf{x}_{1:k}), \quad (12)$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ are two unconstrained functions $\mathbb{R}^{d-k} \rightarrow \mathbb{R}^{d-k}$. The same authors introduced RealNVP [Dinh et al., 2017], which combines coupling layers with normalizing flows and multi-scale architectures for image generation. Glow [Kingma and Dhariwal, 2018] extends RealNVP by introducing invertible 1x1 convolutions between each step of the flow. In this work we have used UMNNs in the context of autoregressive architectures, however UMNNs could also be applied to replace the linear transformation in coupling layers.

Finally, our architecture also shares a connection with Neural Ordinary Differential Equations [NODE, Chen et al., 2018]. The core idea of this architecture is to learn an ordinary differential equation which dynamic is parameterized by a neural network. Training can be carried




Figure 2: Density estimation and sampling with a UMNN-MAF network on 2D toy problems. **Top:** Samples from the empirical distribution $p(\mathbf{x})$. **Middle:** Learned density $p(\mathbf{x}; \theta)$. **Bottom:** Samples drawn by numerical inversion. UMNN-MAF manages to precisely capture multi-modal and/or discontinuous distributions. Sampling is possible even if the model is not invertible analytically.

out by backpropagating efficiently through the ODE solver, with constant memory requirements. Among other applications, NODE can be used to model a continuous normalizing flow with a free-form Jacobian as in FFJORD [Grathwohl et al., 2018]. Similarly, a UMNN transformation can be seen as a structured neural ordinary differential equation in which the dynamic of the vector field is separable and can be solved efficiently by direct integration.

5 Experiments

In this section, we evaluate the expressiveness of UMNN-MAF on a variety of density estimation benchmarks, as well as for approximate inference in variational auto-encoders. The source code to reproduce our experiments will be made available on Github at the end of the reviewing process.

Experiments were carried out using the same integrand neural network in the UMNN component – i.e., in Equation 8, $f^i = f$ with shared weights $\psi^i = \psi$ for $i \in \{1, \dots, d\}$. The functions β^i are taken to be equal to one of the outputs of the embedding network. We observed in our experiments that sharing the same integrand function does not impact performance. Therefore, the neural embedding function \mathbf{h}^i must produce a fixed size output for $i \in \{1, \dots, d\}$.

5.1 2D toy problems

We first train a UMNN-MAF on 2-dimensional toy distributions, as defined by Grathwohl et al. [2018]. To train the model, we minimize the negative log-likelihood of observed data

$$L(\theta) = - \sum_{n=1}^N \left[\log p_Z(\mathbf{g}(\mathbf{x}^n; \theta)) + \sum_{i=1}^d \log f(x_i^n, \mathbf{h}^i(\mathbf{x}_{1:i-1}^n)) \right]. \quad (13)$$

The flow used to solve these tasks is the same for all distributions and is composed of a single transformation. More details can be found in Appendix A.1.

Figure 2 demonstrates that our model is able to learn a change of variables that warps a simple isotropic Gaussian into multimodal and/or discontinuous distributions. We observe from the figure that our model precisely captures the density of the data. We also observe that numerical inversion for generating samples yields good results.

5.2 Density estimation

We further validate UMNN-MAF by comparing it to state-of-the-art normalizing flows. We carry out experiments on tabular datasets (POWER, GAS, HEPMASS, MINIBOONE, BSDS300) as well as on MNIST. We follow the experimental protocol of Papamakarios et al. [2017]. All training hyper-parameters and architectural details are given in Appendix A.1. For each dataset, we report

Table 1: Average negative log-likelihood on test data over 3 runs, error bars are equal to the standard deviation. Results are reported in nats for tabular data and bits/dim for MNIST; lower is better. The best performing architecture for each dataset is written in bold and the best performing architecture per category is underlined. (a) Non-autoregressive models, (b) Autoregressive models, (c) Monotonic and autoregressive models. UMNN outperforms other monotonic transformations on 4 tasks over 6 and is the overall best performing model on 2 tasks over 6.

Dataset	POWER	GAS	HEPMASS	MINIBOONE	BSDS300	MNIST
RealNVP - Dinh et al. [2017]	$-0.17 \pm .01$	$-8.33 \pm .14$	$18.71 \pm .02$	$13.55 \pm .49$	-153.28 ± 1.78	-
(a) Glow - Kingma and Dhariwal [2018]	$-0.17 \pm .01$	$-8.15 \pm .40$	$19.92 \pm .08$	$11.35 \pm .07$	$-155.07 \pm .03$	-
FFJORD - Grathwohl et al. [2018]	<u>$-0.46 \pm .01$</u>	<u>$-8.59 \pm .12$</u>	<u>$14.92 \pm .08$</u>	<u>$10.43 \pm .04$</u>	<u>$-157.40 \pm .19$</u>	-
MADE - Germain et al. [2015]	$3.08 \pm .03$	$-3.56 \pm .04$	$20.98 \pm .02$	$15.59 \pm .50$	$-148.85 \pm .28$	$2.04 \pm .01$
(b) MAF - Papamakarios et al. [2017]	$-0.24 \pm .01$	$-10.08 \pm .02$	$17.70 \pm .02$	$11.75 \pm .44$	$-155.69 \pm .28$	$1.89 \pm .01$
TAN - Oliva et al. [2018]	<u>$-0.60 \pm .01$</u>	<u>$-12.06 \pm .02$</u>	<u>$13.78 \pm .02$</u>	<u>$11.01 \pm .48$</u>	<u>$-159.80 \pm .07$</u>	<u>1.19</u>
NAF - Huang et al. [2018]	$-0.62 \pm .01$	$-11.96 \pm .33$	$15.09 \pm .40$	<u>$8.86 \pm .15$</u>	$-157.73 \pm .30$	-
(c) B-NAF - De Cao et al. [2019]	$-0.61 \pm .01$	<u>$-12.06 \pm .09$</u>	$14.71 \pm .38$	$8.95 \pm .07$	$-157.36 \pm .03$	-
SOS - Jaini et al. [2019]	$-0.60 \pm .01$	$-11.99 \pm .41$	$15.15 \pm .1$	$8.90 \pm .11$	$-157.48 \pm .41$	1.81
UMNN-MAF (ours)	<u>$-0.63 \pm .01$</u>	$-10.89 \pm .7$	<u>$13.99 \pm .21$</u>	$9.67 \pm .13$	<u>$-157.98 \pm .01$</u>	<u>$1.13 \pm .02$</u>

results on test data for our best performing model (selected on the validation data). At testing time we use a large number of integration steps (100) to compute the integral, this ensures its correctness and avoids misestimating the performance of UMNN-MAF.

Table 1 summarizes our results, where we can see that on tabular datasets, our method is competitive with other normalizing flows. For POWER, our architecture slightly outperforms all others. It is also better than other monotonic networks (category (c)) on 3 tabular datasets over 5. From these results, we could conclude that Transformation Autoregressive Networks [TAN, Oliva et al., 2018] is overall the best method for density estimation. It is however important to note that TAN is a flow composed of many heterogeneous transformations (both autoregressive and non-autoregressive). For this reason, it should not be directly compared to the other models which respective results are specific to a single architecture. However, TAN provides the interesting insight that combining heterogeneous components into a flow leads to better results than an homogeneous flow.

Notably, we do not make use of a multi-scale architecture to train our model on MNIST. On this task, UMNN-MAF slightly outperforms all other models by a reasonable margin. Samples generated by a conditional model are shown on Figure 3, for which it is worth noting that UMNN-MAF is the first monotonic architecture that has been inverted to generate samples. Indeed, MNIST can be considered as a high dimensional dataset ($d = 784$) for standard feed forward neural networks which autoregressive networks are part of. NAF and B-NAF do not report any result for this benchmark, presumably because of memory explosion. In comparison, BSDS300, which data dimension is one order of magnitude smaller than MNIST ($63 \ll 784$), are the largest data they have tested on. Table 2 shows the number of parameters used by UMNN-MAF in comparison to B-NAF and NAF. For bigger datasets, UMNN-MAF requires less parameters than NAF to reach similar or better performance. This could explain why NAF has never been used for density estimation on MNIST.




Figure 3: Samples generated by numerical inversion of a conditional UMNN-MAF trained on MNIST. Samples \mathbf{z} are drawn from an isotropic Gaussian with $\sigma = .75$. See Appendix C for more details.

Table 2: Comparison of the number of parameters between NAF, B-NAF and UMNN-MAF. In high dimensional datasets, UMNN-MAF requires fewer parameters than NAF and a similar number to B-NAF.

Dataset	NAF	B-NAF	UMNN-MAF
POWER ($d = 6$)	4.14e5	3.07e5	5.09e5
GAS ($d = 8$)	4.02e5	5.44e5	8.15e5
HEPMASS ($d = 21$)	9.27e6	3.72e6	3.62e6
MINIBOONE ($d = 43$)	7.49e6	4.09e6	3.46e6
BSDS300 ($d = 63$)	3.68e7	8.76e6	1.56e7

Table 3: Average negative evidence lower bound of VAEs over 3 runs, error bars are equal to the standard deviation. Results are reported in bits per dim for Freyfaces and in nats for the other datasets; lower is better. UMNN-NAF is performing slightly better than IAF but is outperformed by B-NAF. We believe that the gap in performance between B-NAF and UMNN is due to the way the NF is conditioned by the encoder’s output.

Dataset	MNIST	Freyfaces	Omniglot	Caltech 101
VAE - Kingma and Welling [2013]	86.65 \pm .06	4.53 \pm .02	104.28 \pm .39	110.80 \pm .46
Planar - Rezende and Mohamed [2015]	86.06 \pm .32	4.40 \pm .06	102.65 \pm .42	109.66 \pm .42
(a) IAF - Kingma et al. [2016]	84.20 \pm .17	4.47 \pm .05	102.41 \pm .04	111.58 \pm .38
Sylvester - Berg et al. [2018]	83.32 \pm .06	4.45 \pm .04	99.00 \pm .04	104.62 \pm .29
FFJORD - Grathwohl et al. [2018]	82.82 \pm .01	4.39 \pm .01	98.33 \pm .09	104.03 \pm .43
B-NAF - De Cao et al. [2019]	83.59 \pm .15	4.42 \pm .05	100.08 \pm .07	105.42 \pm .49
(b) UMNN-MAF (ours)	84.11 \pm .05	4.51 \pm .01	100.98 \pm .13	110.45 \pm .69

5.3 Variational auto-encoders

To assess the performance of our model, we follow the experimental setting of Berg et al. [2018] for VAE. The encoder and the decoder architectures can be found in the appendix of their paper. In VAE it is usual to let the encoder output the parameters of the flow. For UMNN-MAF, this would cause the encoder output’s dimension to be too large. Instead, the encoder output is passed as additional entries of the UMNN-MAF. Like other architectures, the UMNN-MAF also takes as input a vector of noise drawn from an isotropic Gaussian of dimension 64.

Table 3 presents our results. It shows that on MNIST and Omniglot, UMNN-MAF slightly outperforms the classical VAE as well as planar flows. Moreover, on these datasets and Freyfaces, IAF, B-NAF and UMNN-MAF achieve similar results. FFJORD is the best among all, however it is worth noting that the roles of encoder outputs in FFJORD, B-NAF, IAF and Sylvester are all different. We believe that the heterogeneity of the results could be, at least in part, due to the different amortizations.

6 Discussion and summary

Static integral quadrature can be inaccurate. Computing the integral with static Clenshaw-Curtis quadrature only requires the evaluation of the integrand at predefined points. As such, batches of points can be processed all at once, which makes static Clenshaw-Curtis quadrature well suited for neural networks. However, static quadratures do not account for the error made during the integration. As a consequence, the quadrature is inaccurate when the integrand is not smooth enough and the number of integration steps is too small. In this work, we have reduced the integration error by applying the normalization described by Gouk et al. [2018] in order to control the Lipschitz constant of the integrand and appropriately set the number of integration steps. We observed that as long as the Lipschitz constant of the network does not increase dramatically (< 1000), a reasonable number of integration steps (< 100) is sufficient to ensure the convergence of the quadrature. An alternative solution would be to use dynamic quadrature such as dynamic Clenshaw-Curtis.

Efficiency of numerical inversion. Architectures relying on linear transformations [Papamakarios et al., 2017, Kingma et al., 2016, Dinh et al., 2017, Kingma and Dhariwal, 2018] are trivially exactly and efficiently invertible. In contrast, the UMNN transformation has no analytic inverse. Nevertheless, it can be inverted numerically using root-finding algorithms. Since most such algorithms rely on multiple nested evaluations of the function to be inverted, applying them naively to a numerical integral would quickly become very inefficient. However, the Clenshaw-Curtis quadrature is part of the nested quadrature family, meaning that the evaluation of the integral at multiple nested points can take advantage of previous evaluations and thus be implemented efficiently. As an alternative, Oord et al. [2018] have shown that an invertible model can always be distilled to learn its inverse, and thus make the inversion efficient whatever the cost of inversion of the original model.

Scalability and complexity analysis. UMNN-MAF is particularly well suited for density estimation because the computation of the Jacobian only requires a single forward evaluation of a NN.

Together with the Leibniz integral rule, they make the evaluation of the log-likelihood derivative as memory efficient as usual supervised learning, which is equivalent to a single backward pass on the computation graph. By contrast, density estimation with previous monotonic transformations typically requires a backward evaluation of the computation graph of the transformer NN to obtain the Jacobian. Then, this pass must be evaluated backward again in order to obtain the log-likelihood derivative. Both NAF and B-NAF provide a method to make this computation numerically stable, however both fail at not increasing the size of the computation graph of the log-likelihood derivative, hence leading to a memory overhead. The memory saved by the Leibniz rule may serve to speed up the quadrature computation. In the case of static Clenshaw-Curtis, the function values at each evaluation point can be computed in parallel using batch of points. In consequence, when the GPU memory is large enough to store "meta-batches" of size $d \times N \times B$ (with d the dimension of the data, N the number of integration steps and B the batch size) the computation is approximately as fast as a forward evaluation of the integrand network.

Summary We have introduced Unconstrained Monotonic Neural Networks, a new invertible transformation built upon free-form neural networks allowing the use of any state-of-the-art architecture. Monotonicity is guaranteed without imposing constraints on the expressiveness of the hypothesis class, contrary to classical approaches. We have shown that the resulting integrated neural network can be evaluated efficiently using standard quadrature rule while its inverse can be computed using numerical algorithms. We have shown that our transformation can be composed into an autoregressive flow, with competitive or state-of-the-art results on density estimation and variational inference benchmarks. Moreover, UMNN is the first monotonic transformation that has been successfully applied for density estimation on high dimensional data distributions (MNIST), showing better results than the classical approaches.

We identify several avenues for improvement and further research. First, we believe that numerical integration could be fasten up during training, by leveraging the fact that controlled numerical errors can actually help generalization. Moreover, the UMNN transformation would certainly profit from using a dynamic integration scheme, both in terms of accuracy and efficiency. Second, it would be worth comparing the newly introduced monotonic transformation with common approaches for modelling monotonic functions in machine learning. On a similar track, these common approaches could be combined into an autoregressive flow as shown in Section 3.3. Finally, our monotonic transformation could be used within other neural architectures than generative autoregressive networks, such as multi-scale architectures [Dinh et al., 2017] and learnable 1D convolutions [Kingma and Dhariwal, 2018].

Acknowledgments

The authors would like to acknowledge Matthia Sabatelli, Nicolas Vecoven, Antonio Sutera and Louis Wehenkel for useful feedback on the manuscript. They would also like to thank the anonymous reviewers for many relevant remarks. Antoine Wehenkel is a research fellow of the F.R.S.-FNRS (Belgium) and acknowledges its financial support.

References

- N. P. Archer and S. Wang. Application of the back propagation neural network algorithm with monotonicity constraints for two-group classification problems. *Decision Sciences*, 24(1):60–75, 1993.
- R. v. d. Berg, L. Hasenclever, J. M. Tomczak, and M. Welling. Sylvester normalizing flows for variational inference. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pages 6571–6583, 2018.
- H. Daniels and M. Velikova. Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6):906–917, 2010.
- N. De Cao, I. Titov, and W. Aziz. Block neural autoregressive flow. *arXiv preprint arXiv:1904.04676*, 2019.
- L. Dinh, D. Krueger, and Y. Bengio. Nice: Non-linear independent components estimation. In *International Conference in Learning Representations workshop track*, 2015.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. In *International Conference in Learning Representations*, 2017.
- M. Germain, K. Gregor, I. Murray, and H. Larochelle. Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pages 881–889, 2015.
- H. Gouk, E. Frank, B. Pfahringer, and M. Cree. Regularisation of neural networks by enforcing lipschitz continuity. *arXiv preprint arXiv:1804.04368*, 2018.
- W. Grathwohl, R. T. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *International Conference on Machine Learning*, 2018.
- M. Gupta, A. Cotter, J. Pfeifer, K. Voevodski, K. Canini, A. Mangylov, W. Moczydlowski, and A. Van Esbroeck. Monotonic calibrated interpolated look-up tables. *The Journal of Machine Learning Research*, 17(1):3790–3836, 2016.
- D. Ha, A. M. Dai, and Q. V. Le. Hypernetworks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville. Neural autoregressive flows. In *International Conference on Machine Learning*, pages 2083–2092, 2018.
- P. Jaini, K. A. Selby, and Y. Yu. Sum-of-squares polynomial flow. *arXiv preprint arXiv:1905.02325*, 2019.
- D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10236–10245, 2018.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations (ICLR)*, 2013.
- D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016.
- J. Oliva, A. Dubey, M. Zaheer, B. Póczos, R. Salakhutdinov, E. Xing, and J. Schneider. Transformation autoregressive networks. In *International Conference on Machine Learning*, pages 3895–3904, 2018.
- A. Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. Driessche, E. Lockhart, L. Cobo, F. Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International Conference on Machine Learning*, pages 3915–3923, 2018.

- G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.
- G. Papamakarios, D. C. Sterratt, and I. Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538, 2015.
- J. Sill. Monotonic networks. In *Advances in neural information processing systems*, pages 661–667, 1998.
- D. Tran, M. D. Hoffman, R. A. Saurous, E. Brevdo, K. Murphy, and D. M. Blei. Deep probabilistic programming. In *5th International Conference on Learning Representations (ICLR)*, 2017.
- A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. Wavenet: A generative model for raw audio. In *9th ISCA Speech Synthesis Workshop*, pages 125–125, 2016.
- S. You, D. Ding, K. Canini, J. Pfeifer, and M. Gupta. Deep lattice networks and partial monotonic functions. In *Advances in Neural Information Processing Systems*, pages 2981–2989, 2017.

A Experimental setup

A.1 Density estimation and toy problems hyperparameters

Table 4 reports the training configurations for the 2D toy problems and the 5 tabular datasets. For tabular data the best performing architecture has been found after some preliminary experiments, while this was not needed for the 2D toy problems. During our preliminary experiments we tested different integrand network architectures, we tested on the number of hidden layers $L \in \{3, 4\}$ and on their dimension $D \in \{50, 100, 150, 200\}$. The architecture of the embedding networks is the best performing MADE network used in NAF [Huang et al., 2018]. We used the Adam optimizer and tried different learning rate $\lambda \in \{10^{-3}, 5 \times 10^{-4}, 10^{-4}\}$. When the learning rate chosen was greater than 10^{-4} we schedule once the learning rate to 10^{-4} after the first plateau. We also tested for different weights decay values $W \in \{10^{-5}, 10^{-2}\}$. The batch size was chosen to be as big as possible while not harming the learning procedure. We observed during our preliminary experiments that choosing the number of integration steps at random (uniformly from 20 to 100) for each batch regularizes the complexity of the integral. For MNIST, we observed that 25 integration steps was enough if the Lipschitz constant of the network is constraint (with the normalization proposed by Gouk et al. [2018]) to be smaller than 1.5.

Dataset	POWER	GAS	HEPMASS	MINIBOONE	BSDS300	MNIST	2D Toys
Lipschitz	-	-	-	-	2.5	1.5	-
N°integ. steps	rand	rand	rand	rand	rand	25	50
Embedding net	2×100	2×100	2×512	1×512	2×1024	1×1024	4×50
Integrand net ($L \times D$)	4×150	3×200	4×200	3×50	4×150	3×150	4×50
Learning rate (λ)	10^{-3}	10^{-3}	10^{-3}	10^{-3}	10^{-4}	10^{-3}	10^{-3}
N°flows	5	10	5	3	5	5	1
Embedding Size	30	30	30	30	30	30	10
Weight decay (W)	10^{-5}	10^{-2}	10^{-4}	10^{-2}	10^{-2}	10^{-2}	10^{-5}
Batch size	10000	10000	100	500	100	100	100

Table 4: Training configurations for density estimation and toy problems.

A.2 Variational auto-encoders

Table 5 presents the architectural settings of the normalizing flows used inside the variational auto-encoders. The number of values outputted by the encoder is always taken to be equal to 320. These values as well as the 64-dimensional noise vector are the inputs of the embedding network which is constantly made of one hidden layer of 1280 neurons. We have performed a small grid search on the integrand network architecture, we took a look at 2 different number $L \in \{3, 4\}$ of hidden layers of dimensions $D \in \{100, 150\}$.

Dataset	MNIST	Freyfaces	Omniglot	Caltech 101
Lipschitz	-	-	-	-
N°integ. steps	rand	rand	rand	rand
Encoder Output	320	320	320	320
Embedding net	1×1280	1×1280	1×1280	1×1280
Integrand net	4×100	3×100	4×100	4×100
N°flows	16	8	16	16
Embedding Size	30	30	30	30

Table 5: Training configurations of variational auto-encoder.

B Clenshaw-Curtis module

Algorithm 1 Clenshaw-Curtis quadrature

Input: x : A tensor of scalar values that represent the superior integration bounds.
 \mathbf{h} : A tensor of vectors that representing embeddings.

Output: F : A tensor of scalar values that represent the integral of $\int_0^x f(t; \mathbf{h}) dt$.

Hyper-parameters: f : A derivable function $\mathbb{R} \rightarrow \mathbb{R}$.
 N : The number of integration steps.

- 1: **procedure** FORWARD($x, \mathbf{h}; f, N$)
- 2: \triangleright Compute weights and evaluation steps for Clenshaw-Curtis quadrature
- 3: $\mathbf{w}, \delta_x = \text{COMPUTE_CC_WEIGHTS}(N)$
- 4: $F = 0$
- 5: **for** $i \in [1, N]$ **do**
- 6: $x_i = x_0 + \frac{1}{2}(x - x_0)(\delta_x[i] + 1)$ \triangleright Compute the next point to evaluate
- 7: $\delta_F = f(x_i; \mathbf{h})$
- 8: $F = F + \mathbf{w}[i]\delta_F$
- 9: **end for**
- 10: $F = \frac{F}{2}(x - x_0)$
- 11: **return** F
- 12: **end procedure**

Inputs: x : A tensor of scalar values that represent the superior integration bounds.
 \mathbf{h} : A tensor of vectors that representing embeddings.
 ∇_{out} : The derivatives of the loss function with respect to $\int_0^x f(t; \mathbf{h}) dt$ for all x .

Outputs: ∇_x : The gradient of $\int_0^x f(t; \mathbf{h}) dt$ with respect to x .
 ∇_{θ} : The gradient of $\int_0^x f(t; \mathbf{h}) dt$ with respect to \mathbf{f} parameters.
 $\nabla_{\mathbf{h}}$: The gradient of $\int_0^x f(t; \mathbf{h}) dt$ with respect to \mathbf{h} .

Hyper-parameters: f : A derivable function $\mathbb{R} \rightarrow \mathbb{R}$.
 N : The number of integration steps.

- 1: **procedure** BACKWARD($x, \mathbf{h}, \nabla_{out}; f, N$)
- 2: \triangleright Compute weights and evaluation steps for Clenshaw-Curtis quadrature
- 3: $\mathbf{w}, \delta_x = \text{COMPUTE_CC_WEIGHTS}(N)$
- 4: $F, \nabla_{\theta}, \nabla_{\mathbf{h}} = 0, 0, 0$
- 5: **for** $i \in [1, N]$ **do**
- 6: $x_i = x_0 + \frac{1}{2}(x - x_0)(\delta_x[i] + 1)$ \triangleright Compute the next point to evaluate
- 7: $\delta_F = f(x_i; \mathbf{h})$
- 8: \triangleright Sum up for all samples of the batch the gradients with respect to inputs \mathbf{h}
- 9: $\delta_{\nabla_{\mathbf{h}}} = \sum_{j=1}^B \nabla_{\mathbf{h}^j} (\delta_F^j) \nabla_{out}^j (x^j - x_0^j)$
- 10: \triangleright Sum up for all samples of the batch the gradients with respect to parameters θ
- 11: $\delta_{\nabla_{\theta}} = \sum_{j=1}^B \nabla_{\theta} (\delta_F^j) \nabla_{out}^j (x^j - x_0^j)$
- 12: $\nabla_{\mathbf{h}} = \nabla_{\mathbf{h}} + \mathbf{w}[i]\delta_{\nabla_{\mathbf{h}}}$
- 13: $\nabla_{\theta} = \nabla_{\theta} + \mathbf{w}[i]\delta_{\nabla_{\theta}}$
- 14: **end for**
- 15: \triangleright Gradients with respect to superior integration bound.
- 16: $\nabla_x = f(x, \mathbf{h})\nabla_{out}$
- 17: **return** $\nabla_x, \nabla_{\theta}, \nabla_{\mathbf{h}}$
- 18: **end procedure**

C Generated images from MNIST

Figure 4 presents samples generated from two UMNN-MAF trained on MNIST, respectively with (sub-figure a) and without (sub-figure b) labels. The samples are generated with different levels of noise, which are the product of the inversion of the network with random values drawn from $\mathcal{N}(0, T)$, with T being the sampling temperature. The sampling temperature increases linearly from 0.1 (top rows) to 1.0 (bottom rows). We can observe that the unconditional model fails to incorporate digit structure when the level of noise is too small. However, when the level is sufficient it is able to generate random digits with a high level of heterogeneity.




Figure 4: (a): Class-conditional generated images from MNIST. The temperature of sampling increases from 0.1 (top row) to 1.0 (bottom row). Columns correspond to different classes. (b): Unconditional generated images from MNIST. The temperature of sampling goes from 0.1 at top row to 1.0 at bottom row. Columns are different random noise values.