

What?


UMNN is a new neural network architecture for modelling monotonic functions.

How?

The strictly positive scalar output of a neural network is numerically integrated.

Applications?

We combine UMNNs with autoregressive flows to perform density estimation.



Code




Arxiv

1908.05164

Monotonic Networks (UMNN)


A necessary and sufficient condition for a continuously derivable function to be strictly monotonic is for its derivative to be of constant sign.



Monotonicity is hard to enforce

BUT

positiveness is not!




Derivation



Solution:
Model the derivative and integrate

$$\hat{y} = F(x; \theta) = \int_0^x \frac{\partial f(t; \theta)}{\partial t} dt + \text{const}$$

Forward: Integration



Update

$$\mathcal{L}_\theta = \sum_i l(\hat{y}_i, y_i)$$

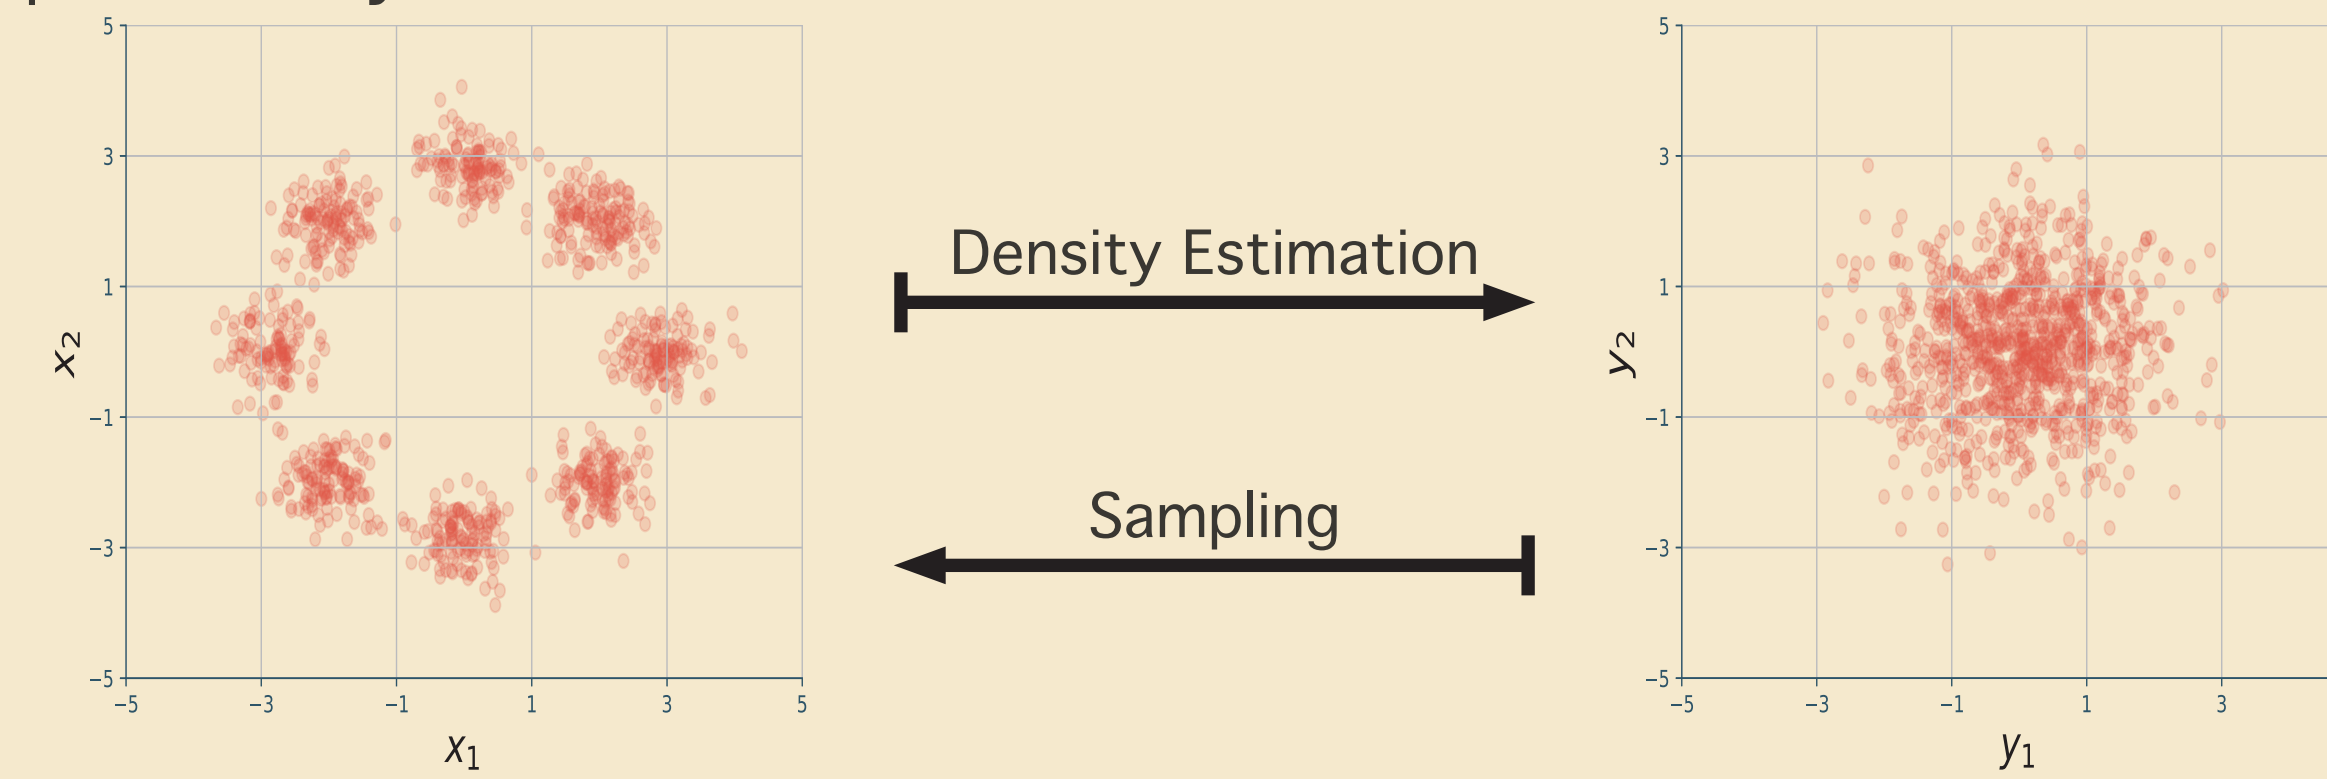
Backward: Integration

$$\frac{d\mathcal{L}_\theta}{d\theta} = \sum_i \frac{\partial l(\hat{y}_i, y_i)}{\partial \hat{y}_i} \frac{d\hat{y}_i}{d\theta}$$

$$\frac{d\hat{y}_i}{d\theta} = \frac{dF(x; \theta)}{d\theta} = \int_0^x \frac{d(\frac{\partial f(t; \theta)}{\partial t})}{d\theta} dt$$

Change of variables

The adequate bijective function combined with any base distribution is able to represent any continuous random variable.



$$p(\mathbf{x}; \theta) = p_{\mathbf{y}}(\mathbf{y} = g(\mathbf{x}; \theta)) |\det J_g(\mathbf{x}; \theta)|, \quad g(\cdot, \theta) \text{ an invertible neural network.}$$

Autoregressive Density Estimation

Autoregressive transformations can be written as

$$g(\mathbf{x}; \theta) = [g_1(x_1; \theta) \dots g_i(\mathbf{x}_{1:i}; \theta) \dots g_d(\mathbf{x}_{1:d}; \theta)]$$

Invertible? Yes, if each scalar transformation in h^i is itself invertible.

$$x_1 = g_1^{-1}(y_1; \theta)$$

$$x_i = g_i^{-1}(y_i, \mathbf{x}_{1:i-1}; \theta)$$

The induced multivariate density can be expressed by the chain rule as

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

We build UMNN into an autoregressive transformation as

$$g^i(\mathbf{x}_{1:i}; \theta) = F^i(x_i, \mathbf{h}^i(\mathbf{x}_{1:i-1})) = \int_0^{x_i} f^i(t, \mathbf{h}^i(\mathbf{x}_{1:i-1})) dt + \beta^i(\mathbf{h}^i(\mathbf{x}_{1:i-1}))$$

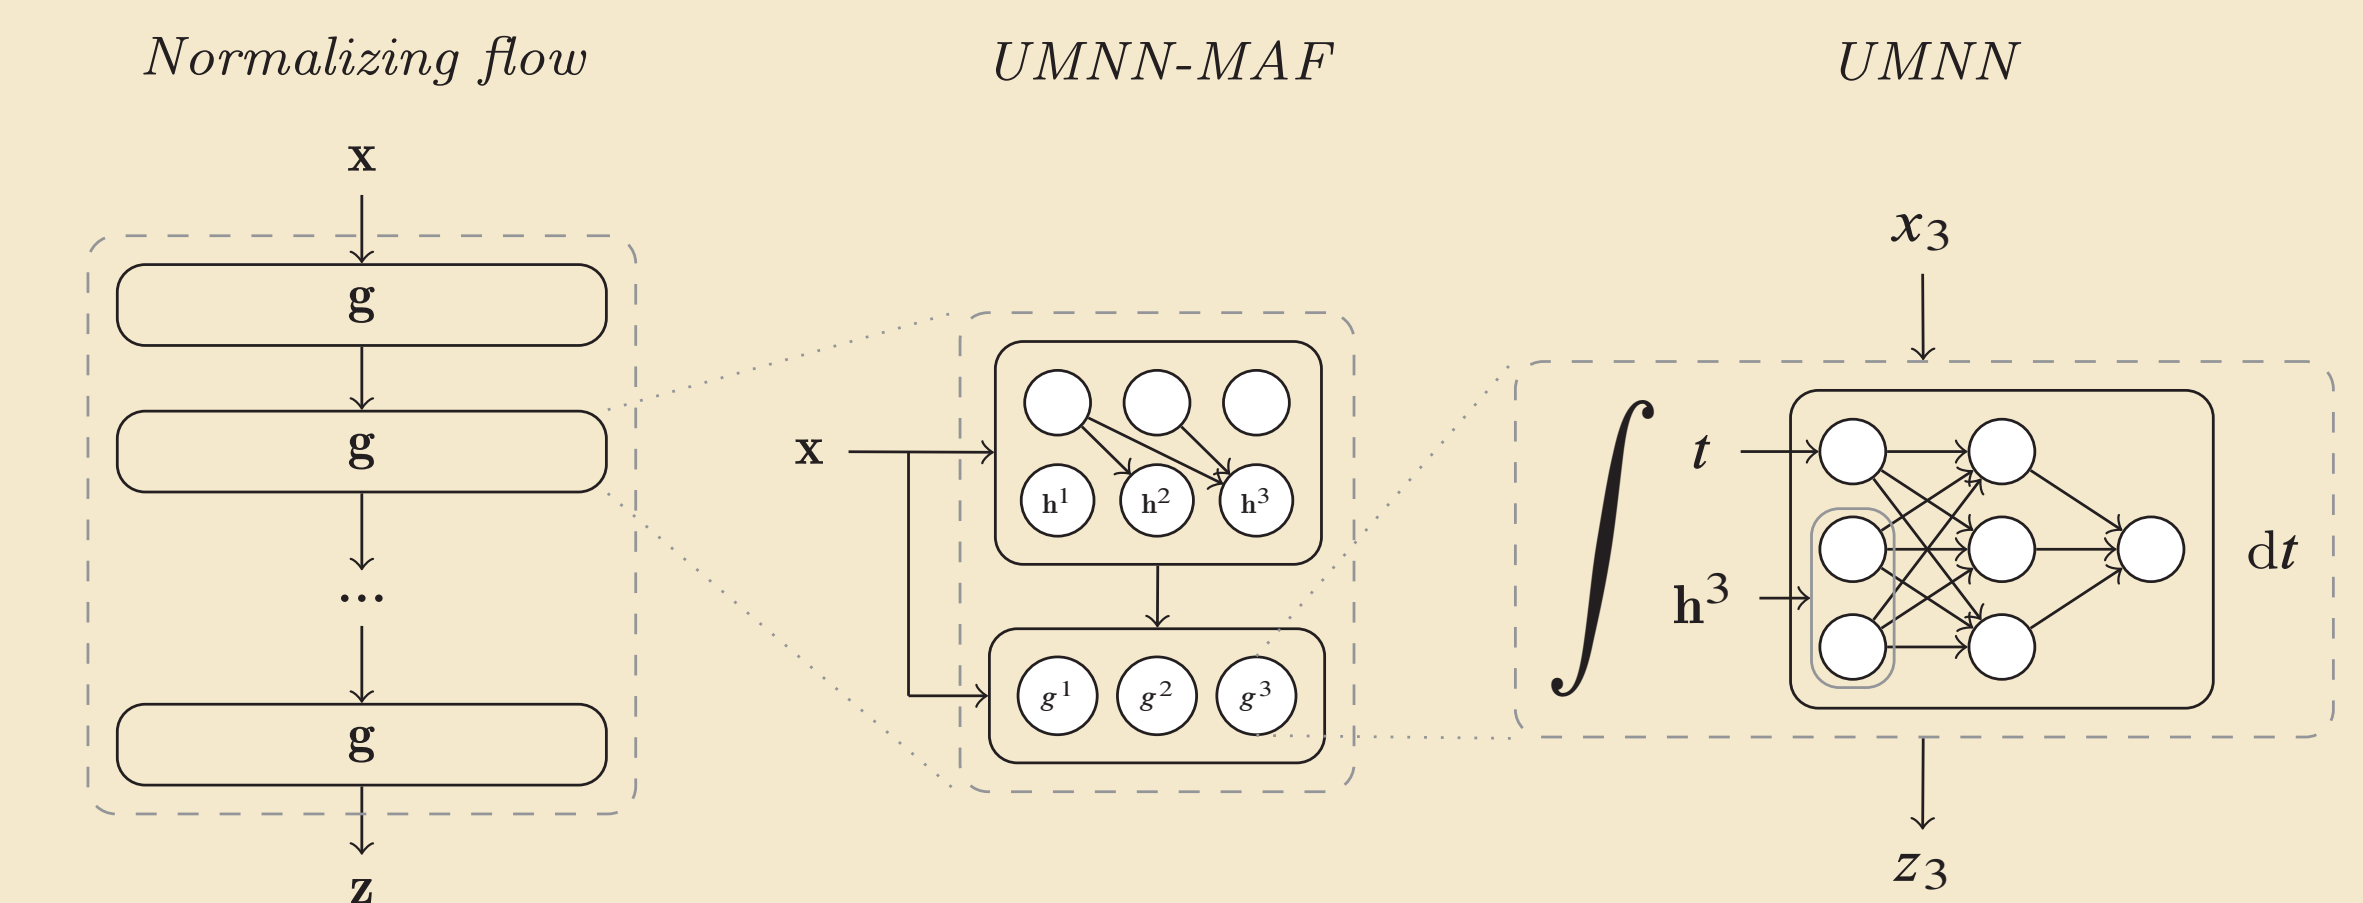
The autoregressive embeddings \mathbf{h}^i are computed with a masked autoregressive network.

We call this transformation UMNN-MAF, it leads to the following elegant expression of the log-likelihood of a point given the model:

$$\log p(\mathbf{x}; \theta) = \log p_{\mathbf{y}}(g(\mathbf{x}; \theta)) + \sum_{i=1}^d \log f^i(x_i, \mathbf{h}^i(\mathbf{x}_{1:i-1})).$$


Architecture

The combination of a UMNN with autoregressive transformations can be visualized as follows:



Results

Toy problems



Density estimation

Dataset	POWER	GAS	HEPMAS	MINIBOONE	BSDS300	MNIST
RealNVP - Dinh et al. [2017]	-0.17 ± 0.01	-8.33 ± 1.14	18.71 ± 0.02	13.55 ± 0.49	-153.28 ± 1.74	-
(a) Glow - Kingma and Dhariwal [2018]	-0.17 ± 0.01	-8.15 ± 0.40	19.92 ± 0.08	11.35 ± 0.07	-155.07 ± 0.03	-
FFJORD - Grathwohl et al. [2018]	-0.46 ± 0.01	-8.59 ± 0.12	14.92 ± 0.08	10.43 ± 0.04	-157.40 ± 0.10	-
MADE - Germain et al. [2015]	3.08 ± 0.03	-3.56 ± 0.04	20.98 ± 0.02	15.59 ± 0.50	-148.85 ± 0.28	2.04 ± 0.01
(b) MAF - Papamakarios et al. [2017]	-0.24 ± 0.01	-10.08 ± 0.02	17.70 ± 0.02	11.75 ± 0.44	-155.69 ± 0.28	1.89 ± 0.01
TAN - Oliva et al. [2018]	-0.60 ± 0.01	-12.00 ± 0.02	13.78 ± 0.02	11.01 ± 0.48	-159.80 ± 0.07	1.19 ± 0.01
NAF - Huang et al. [2018]	-0.62 ± 0.01	-11.96 ± 0.33	15.09 ± 0.40	8.80 ± 0.15	-157.73 ± 0.30	-
(c) B-NAF - De Cao et al. [2019]	-0.61 ± 0.01	-12.06 ± 0.09	14.71 ± 0.38	8.95 ± 0.07	-157.36 ± 0.03	-
SOS - Jaini et al. [2019]	-0.60 ± 0.01	-11.99 ± 0.41	15.15 ± 0.11	8.90 ± 0.11	-157.48 ± 0.41	1.81 ± 0.01
UMNN-MAF (ours)	-0.63 ± 0.01	-10.89 ± 0.7	13.99 ± 0.21	9.67 ± 0.13	-157.98 ± 0.01	1.13 ± 0.02

MNIST



Take home messages

- Any monotonic function can be modeled by a neural network that represents the function derivative.
- The backward pass is memory efficient thanks to Leibniz rule.
- UMNNs can be used as building blocks of autoregressive bijective maps and provide a state of the art density estimator.