
You say Normalizing Flows I see Bayesian Networks

Antoine Wehenkel¹ Gilles Louppe¹

Abstract

Normalizing flows have emerged as an important family of deep neural networks for modelling complex probability distributions. In this note, we revisit their coupling and autoregressive transformation layers as probabilistic graphical models and show that they reduce to Bayesian networks with a pre-defined topology and a learnable density at each node. From this new perspective, we provide three results. First, we show that stacking multiple transformations in a normalizing flow relaxes independence assumptions and entangles the model distribution. Second, we show that a fundamental leap of capacity emerges when the depth of affine flows exceeds 3 transformation layers. Third, we prove the non-universality of the affine normalizing flow, regardless of its depth.

1. Introduction

Normalizing flows (NF, 5) have gained popularity in the recent years because of their unique ability to model complex data distributions while allowing both for sampling and exact density computation. This family of deep neural networks combines a base distribution with a series of invertible transformations while keeping track of the change of density that is caused by each transformation.

Probabilistic graphical models (PGMs) are well-established mathematical tools that combine graph and probability theory to ease the manipulation of joint distributions. They are commonly used to visualize and reason about the set of independencies in probabilistic models. Among PGMs, Bayesian networks (BN, 4) offer a nice balance between readability and modeling capacity. Reading independencies stated by a BN is simple and can be performed graphically with the d-separation algorithm (2).

In this note, we revisit NFs as Bayesian networks. We

¹University of Liege. Correspondence to: Antoine Wehenkel <antoine.wehenkel@uliege.be>.

Second workshop on *Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models* (ICML 2020), Virtual Conference

first briefly review the mathematical grounds of these two worlds. Then, for the first time in the literature, we show that the modeling assumptions behind coupling and autoregressive transformations can be perfectly expressed by distinct classes of BNs. From this insight, we show that stacking multiple transformation layers relaxes independencies and entangles the model distribution. Then, we show that a fundamental change of regime emerges when the NF architecture includes 3 transformation steps or more. Finally, we prove the non-universality of affine normalizing flows.

2. Background

2.1. Normalizing flows

A normalizing flow is defined as a sequence of invertible transformation steps $g_k : \mathbb{R}^d \rightarrow \mathbb{R}^d$ ($k = 1, \dots, K$) that are composed together to create an expressive invertible mapping $\mathbf{g} = g_1 \circ \dots \circ g_K : \mathbb{R}^d \rightarrow \mathbb{R}^d$. This mapping can be used to perform density estimation, using $\mathbf{g}(\cdot; \theta) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ to map a sample $\mathbf{x} \in \mathbb{R}^d$ to a latent vector $\mathbf{z} \in \mathbb{R}^d$ equipped with a density $p_{\mathbf{z}}(\mathbf{z})$. The transformation \mathbf{g} implicitly defines a density $p(\mathbf{x}; \theta)$ as given by the change of variables formula,

$$p(\mathbf{x}; \theta) = p_{\mathbf{z}}(\mathbf{g}(\mathbf{x}; \theta)) \left| \det J_{\mathbf{g}(\mathbf{x}; \theta)} \right|,$$

where $J_{\mathbf{g}(\mathbf{x}; \theta)}$ is the Jacobian of $\mathbf{g}(\mathbf{x}; \theta)$ with respect to \mathbf{x} . The resulting model is trained by maximizing the likelihood of the data $\{\mathbf{x}^1, \dots, \mathbf{x}^N\}$. NFs can also be used for data generation tasks while keeping track of the density of the generated samples such as to improve the latent distribution in variational auto-encoders (5). In the rest of this paper, we will not distinguish between \mathbf{g} and g_k when the discussion will be focused on only one of these steps g_k .

In general, steps \mathbf{g} can take any form as long as they define a bijective map. Here, we focus on a sub-class of normalizing flows for which these steps can be mathematically described as

$$\mathbf{g}(\mathbf{x}) = \begin{bmatrix} g^1(x_1; \mathbf{c}^1(\mathbf{x})) & \dots & g^d(x_d; \mathbf{c}^d(\mathbf{x})) \end{bmatrix},$$

where the \mathbf{c}^i are denoted as the **conditioners** and constrain the structure of the Jacobian of \mathbf{g} . The functions g^i , referred as **normalizers**, are partially parameterized by their conditioner and must be invertible with respect to their input

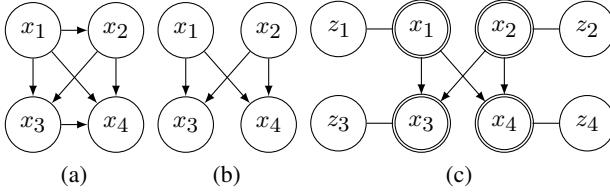


Figure 1. Bayesian networks for single-step normalizing flows on a vector $\mathbf{x} = [x_1, x_2, x_3, x_4]^T$. (a) BN for an autoregressive conditioner. (b) BN for a coupling conditioner. (c) Pseudo BN for a coupling conditioner, with the latent variables shown explicitly. Double circles stand for deterministic functions of the parents and non-directed edges stand for bijective relationships.

variable x_i . These are usually defined as affine or strictly monotonic functions, with the latter being the most general class of invertible scalar continuous functions. In this note, we mainly discuss affine normalizers that can be expressed as $g(x; m, s) = x \exp(s) + m$ where $m \in \mathbb{R}$ and $s \in \mathbb{R}$ are computed by the conditioner.

2.2. Bayesian networks

Bayesian networks allow for a compact and natural representation of probability distributions by exploiting conditional independence. More precisely, a BN is a directed acyclic graph (DAG) which structure encodes for the conditional independencies through the concept of d-separation (2). Equivalently, its skeleton supports an efficient factorization of the joint distribution.

A BN is able to model a distribution p if and only if it is an I-map with respect to p . That is, iff the set of independencies stated by the BN structure is a subset of the independencies that holds for p . Equivalently, a BN is a valid representation of a random vector \mathbf{x} iff its density $p_{\mathbf{x}}(\mathbf{x})$ can be factorized by the BN structure as

$$p_{\mathbf{x}}(\mathbf{x}) = \prod_{i=1}^d p(x_i | \mathcal{P}_i), \quad (1)$$

where $\mathcal{P}_i = \{j : A_{i,j} = 1\}$ denotes the set of parents of the vertex i and $A \in \{0, 1\}^{d \times d}$ is the adjacency matrix of the BN. As an example, Fig. 1a is a valid BN for any distribution over \mathbf{x} because it does not state any independence, leading to a factorization that results in the chain rule.

3. Normalizing flows as Bayesian networks

3.1. Autoregressive conditioners

Autoregressive conditioners can be expressed as

$$c^i(\mathbf{x}) = \mathbf{h}^i \left(\begin{bmatrix} x_1 & \dots & x_{i-1} \end{bmatrix}^T \right),$$

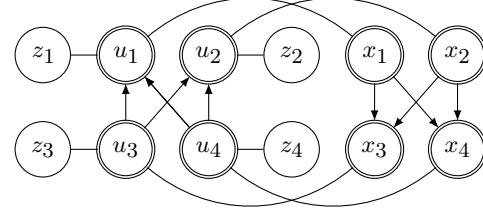


Figure 2. A Bayesian network equivalent to a 2-step normalizing flow based on coupling layers. Independence statements are relaxed by the second step.

where $\mathbf{h}^i : \mathbb{R}^{i-1} \rightarrow \mathbb{R}^l$ are functions of the first $i-1$ components of \mathbf{x} and whose output size depends on architectural choices. These conditioners constrain the Jacobian of g to be lower triangular, making the computation of its determinant $\mathcal{O}(d)$. The multivariate density $p(\mathbf{x}; \theta)$ induced by $\mathbf{g}(\mathbf{x}; \theta)$ and $p_{\mathbf{z}}(\mathbf{z})$ can be expressed as a product of d univariate conditional densities,

$$p(\mathbf{x}; \theta) = p(x_1; \theta) \prod_{i=2}^d p(x_i | \mathbf{x}_{1:i-1}; \theta). \quad (2)$$

When $p_{\mathbf{z}}(\mathbf{z})$ is a factored distribution $p_{\mathbf{z}}(\mathbf{z}) = \prod_{i=1}^d p(z_i)$, we identify that each component z_i coupled with the corresponding function g^i encodes for the conditional $p(x_i | \mathbf{x}_{1:i-1}; \theta)$. An explicit connection between BNs and autoregressive conditioners can be made if we define $\mathcal{P}_i = \{x_1, \dots, x_{i-1}\}$ and compare (2) with (1). Therefore, and as illustrated in Fig. 1a, autoregressive conditioners can be seen as a way to model the conditional factors of a BN that does not state any independence.

3.2. Coupling conditioners

Coupling conditioners (1) are another popular type of conditioners used in normalizing flows. The conditioners c^i made from coupling layers are defined as

$$c^i(\mathbf{x}) = \begin{cases} \underline{\mathbf{h}}^i & \text{if } i < k \\ \mathbf{h}^i(\mathbf{x}_{<k}) & \text{if } i \geq k \end{cases}$$

where the $\underline{\mathbf{h}}^i$ symbol define constant values. As for autoregressive conditioners, the Jacobian of g made of coupling layers is lower triangular. Assuming a factored latent distribution, the density associated with these conditioners can be written as follows:

$$p(\mathbf{x}; \theta) = \prod_{i=1}^{k-1} p(x_i) \prod_{i=k}^d p(x_i | \mathbf{x}_{<k}),$$

$$\text{where } p(x_i) = p(g^i(x_i; \underline{\mathbf{h}}^i)) \frac{\partial g^i(x_i; \underline{\mathbf{h}}^i)}{\partial x_i}$$

$$\text{and } p(x_i | \mathbf{x}_{<k}) = p(g^i(x_i; \mathbf{h}^i(\mathbf{x}_{<k}))) \frac{\partial g^i(x_i; \mathbf{h}^i(\mathbf{x}_{<k}))}{\partial x_i}.$$

The factors define valid 1D conditional probability distributions because they can be seen as 1D changes of variables between z_i and x_i . This factorization can be graphically expressed by a BN as shown in Fig. 1b. In addition, we can see Fig. 1b as the marginal BN of Fig. 1c which fully describes the stochastic process modeled by a NF that is made of a single transformation step and a coupling conditioner. In contrast to autoregressive conditioners, coupling layers are not by themselves universal density approximators, even when associated with very expressive normalizers g^i . Indeed, d-separation reveals independencies stated by this class of BN, such as the conditional independence between each pair in $\mathbf{x}_{\geq k}$ knowing $\mathbf{x}_{< k}$. These independence statements do not hold in general.

3.3. Stacking transformation steps

In practice, the invertible transformations discussed above are often stacked together in order to increase the representation capacity of the flow, with the popular good practice of permuting the vector components between two transformation steps. The structural benefits of this stacking strategy can be explained from the perspective of the underlying BN.

First, a BN that explicitly includes latent variables is faithful as long as the sub-graph made only of those latent nodes is an I-map with respect to their distribution. Normalizing flows composed of multiple transformation layers can therefore be viewed as single transformation flows whose latent distribution is itself recursively modeled by a normalizing flow. As an example, Fig. 2 illustrates a NF made of two transformation steps with coupling conditioners. It can be observed that the latent vector \mathbf{u} is itself a normalizing flow whose distribution can be factored out by a class of BN.

Second, from the BN associated to a NF, we observe that additional layers relax the independence assumptions defined by its conditioners. The distribution modeled by the flow gets more entangled at each additional layer. For example, Fig. 2 shows that for coupling layers, the additional steps relax the strong conditional independencies between x_1 and x_2 of the single transformation NF of Fig. 1c. Indeed, we can observe from the figure that x_1 and x_2 have common ancestors (z_3 and z_4) whereas they are clearly assumed independent in Fig. 1b.

In general, we note that edges between two nodes in a BN do not model dependence, only the absence of edges does model independence. However, because some of the relationship between nodes are bijective, this implies that these nodes are strictly dependent on each other. We represent these relationships with undirected edges in the BN, as it can be seen in Fig. 2.

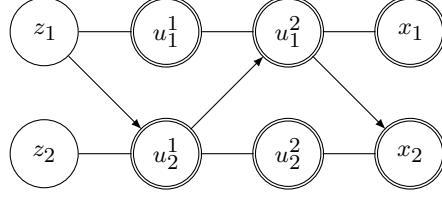


Figure 3. The Bayesian network of a three-steps normalizing flow on vector $\mathbf{x} = [x_1, x_2]^T \in \mathbb{R}^4$. It can be observed that the distribution of the intermediate latent variables, and at the end of the vector \mathbf{x} , becomes more entangled at each additional transformation step. Considering the undirected edges as affine relationships, we see that while u_1^1 and u_2^1 are affine transformations of the latent \mathbf{z} , the vector \mathbf{x} cannot be expressed as a linear function of the latent \mathbf{z} .

4. Affine normalizing flows unlock their capacity with 3 transformation steps

We now show how some of the limitations of affine normalizers can be relaxed by stacking multiple transformation steps. We also discuss why some limitations cannot be relaxed even with a large number of transformation steps. We intentionally put aside monotonic normalizers because they have already been proven to lead to universal density approximators when the conditioner is autoregressive (3). We focus our discussion on a multivariate normal with an identity covariance matrix as base distribution $p_{\mathbf{z}}(\mathbf{z})$.

We first observe from Fig. 1 that in a NF with a single transformation step at least one component of \mathbf{x} is a function of only one latent variable. If the normalizer is affine and the base distribution is normal, then this necessarily implies that the marginal distribution of this component is normal as well, which will very likely not lead to a good fit. We easily see that adding steps relaxes this constraint. A more interesting question to ask is what exactly the modeling capacity gain for each additional step of affine normalizer is. Shall we add steps to increase capacity or shall we increase the capacity of each step instead? We first discuss a simple 2-dimensional case, which has the advantage of unifying the discussion for autoregressive and coupling conditioners, and then extend it to a more general setting.

Affine NFs made of a single transformation step induce strong constraints on the form of the density. In particular, these models implicitly assume that the data distribution can be factorized as a product of conditional normal distributions. These assumptions are relaxed when accumulating steps in the NF. As an example, Fig. 3 shows the equivalent BN of a 2D NF composed of 3 steps. This flow is mathematically described with the following set of equations:

$$\begin{aligned} u_1^1 &:= z_1 & u_2^1 &:= \exp(s_2^1(z_1))z_2 + m_2^1(z_1) \\ u_2^2 &:= u_2^1 & u_1^2 &:= \exp(s_1^2(u_2^1))z_1 + m_1^2(u_2^1) \\ x_1 &:= u_1^2 & x_2 &:= \exp(s_2^3(u_1^2))u_2^2 + m_2^3(u_1^2) \end{aligned}$$

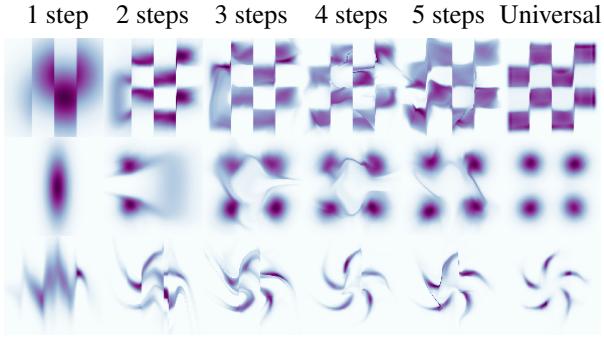


Figure 4. Evolution of an affine normalizing flow’s capacity as the number of steps increases. For comparison, the density learned by a universal density approximator is shown on the last column.

From these equations, we see that after one step the latent variables u_1^1 and u_2^1 are respectively normal and conditionally normal. This is relaxed with the second step, where the latent variable u_1^2 is a non-linear function of two random variables distributed normally (by assumption on the distribution of z_1 and z_2). However, u_2^2 is a stochastic affine transformation of a normal random variable. In addition, we observe that the expression of u_1^2 is strictly more expressive than the expression of u_2^2 . Finally, x_1 and x_2 are non-linear functions of both latent variables z_1 and z_2 . Assuming that the functions s_j^i and m_j^i are universal approximators, we argue that the stochastic process that generates x_1 and the one that generates x_2 are as expressive as each other. Indeed, by making the functions arbitrarily complex the transformation for x_1 could be made arbitrarily close to the transformations for x_2 and vice versa. This is true because both transformations can be seen as an affine transformation of a normal random variables whose scaling and offset factors are non-linear arbitrarily expressive transformations of all the latent variables. Because of this equilibrium between the two expressions, additional steps do not improve the modeling capacity of the flow. The same observations can be made empirically as illustrated in Fig. 3 for 2-dimensional toy problems. A clear leap of capacity occurs from 2-step to 3-step NFs, while having 4 steps or more does not result in any noticeable improvement when s_j^i and m_j^i already have enough capacity.

For $d > 2$, autoregressive and coupling conditioners do not correspond to the same set of equations or BN. However, if the order of the vector is reversed between two transformation steps, the discussion generalizes to any value of d for both conditioners. Indeed, in both cases each component of the intermediate latent vectors can be seen as having a set of conditioning variables and a set of independent variables. At each successive step the indices of the non-conditioning variables are exchanged with the conditioning ones and thus any output vector’s component can be expressed either as a component of the vector form of x_1 or of x_2 .



Figure 5. The equivalent BN of a component with a unique latent variable as ancestor.

5. Affine normalizing flows are not universal density approximators

We argue that affine normalizers do not lead to universal density approximators in general, even for an infinite number of steps. In the following, we assume again that the latent variables are distributed according to a normal distribution with a unit covariance matrix.

To prove the non-universality of affine normalizing flows, one only needs to provide a counter-example. Let us consider the simple setup in which one component x_I of the random vector \mathbf{x} is independent from all the other components. Let us also assume that x_I is distributed under a non-normal distribution. We can then consider two cases. First, x_I has only one component of the latent vector \mathbf{z} as an ancestor. This implies that the equivalent BN would be as in Fig. 5, hence that x_I is a linear function of this ancestor and is therefore normally distributed. Else, x_I has n components of the latent vector as ancestors. However, this second case would imply that at least one undirected edge is removed from the original BN considered in Section 3.3. This cannot happen since it would deadly hurt the bijectivity of the flow.

Besides proving the non-universality of affine NFs, this discussion provides the important insight that when affine normalizers must transform non-linearly some latent variables they introduce dependence in the model of the distribution. In some sense, this means that the additional disorder required to model this non-normal component is performed at the cost of some loss in entropy caused by mutual information between the random vector components.

6. Summary

In this preliminary work, we have revisited normalizing flows from the perspective of Bayesian networks. We have shown that stacking multiple transformations in a normalizing flow relaxes independence assumptions and entangles the model distribution. Then, we have shown that affine normalizing flows benefit from having at least 3 transformation layers. Finally, we demonstrated that they remain non-universal density approximators regardless of their depths.

We hope these results will give practitioners more intuition in the design of normalizing flows. We also believe that this work may lead to further research. First, unifying Bayesian networks and normalizing flows could be pushed one step further with conditioners that are specifically designed to model Bayesian networks. Second, the study could be ex-

tended for other type of normalizing flows such as non-autoregressive monotonic flows. Finally, we believe this study may spark research at the intersection of structural equation modeling, causal networks and normalizing flows.

Acknowledgments

The authors would like to thank Matthia Sabatelli, Johann Brehmer and Louis Wehenkel for proofreading the manuscript. Antoine Wehenkel is a research fellow of the F.R.S.-FNRS (Belgium) and acknowledges its financial support. Gilles Louppe is recipient of the ULiège - NRB Chair on Big data and is thankful for the support of NRB.

References

- [1] L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. In *International Conference in Learning Representations*, 2017.
- [2] D. Geiger, T. Verma, and J. Pearl. d-separation: From theorems to algorithms. In *Machine Intelligence and Pattern Recognition*, volume 10, pages 139–148. Elsevier, 1990.
- [3] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville. Neural autoregressive flows. In *International Conference on Machine Learning*, pages 2083–2092, 2018.
- [4] J. Pearl. Bayesian networks. 2011.
- [5] D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538, 2015.