
Automatic Variational Inference with Cascading Flows

Luca Ambrogioni¹ Gianluigi Silvestri^{1,2} Marcel van Gerven¹

Abstract

The automation of probabilistic reasoning is one of the primary aims of machine learning. Recently, the confluence of variational inference and deep learning has led to powerful and flexible automatic inference methods that can be trained by stochastic gradient descent. In particular, normalizing flows are highly parameterized deep models that can fit arbitrarily complex posterior densities. However, normalizing flows struggle in highly structured probabilistic programs as they need to relearn the forward-pass of the program. Automatic structured variational inference (ASVI) remedies this problem by constructing variational programs that embed the forward-pass. Here, we combine the flexibility of normalizing flows and the prior-embedding property of ASVI in a new family of variational programs, which we named cascading flows. A cascading flows program interposes a newly designed highway flow architecture in between the conditional distributions of the prior program such as to steer it toward the observed data. These programs can be constructed automatically from an input probabilistic program and can also be amortized automatically. We evaluate the performance of the new variational programs in a series of structured inference problems. We find that cascading flows have much higher performance than both normalizing flows and ASVI in a large set of structured inference problems.

1. Introduction

The aim of probabilistic programming is to provide a fully automated software system for statistical inference on arbitrary user-specified probabilistic models (also referred to as probabilistic programs) (Milch et al., 2007; Sato, 1997;

¹Donders Centre for Cognition, Radboud University, Netherlands ²OnePlanet Research Center, imec-the Netherlands, Wageningen, Netherlands. Correspondence to: Luca Ambrogioni <l.ambrogioni@donders.ru.nl>.

Kersting and De Raedt, 2007; Pfeffer, 2001; Park et al., 2005; Goodman et al., 2012; Wingate et al., 2011; Patil et al., 2010; Dillon et al., 2017; Bingham et al., 2019; Tran et al., 2017; 2016). When the probabilistic programs are built from differentiable components, stochastic variational inference (VI) by automatic differentiation offers an effective and computationally efficient solution (Bishop et al., 2002; Kucukelbir et al., 2017; Tran et al., 2017). However, the performance of stochastic VI depends strongly on the choice of a variational program (also known as variational family or variational guide) (Bishop and Winn, 2003).

Until recently, automatic variational program construction was limited to simple mean field and multivariate normal coupling approaches (Kucukelbir et al., 2017). These programs are highly constrained and exploit only minimally the structure of the original probabilistic program. For example, the multivariate normal approach only depends on the number of variables and their support. More recently, normalizing flows revolutionized the field of variational inference by offering a highly flexible parametric model for complex multivariate distributions. Normalizing flows can be easily used to construct automatic variational programs by mapping all the variables to a spherical normal latent space through a learnable diffeomorphism (Rezende and Mohamed, 2015; Papamakarios et al., 2019). While highly flexible, the normalizing flow approach still uses only minimal information concerning the input probabilistic program. This leads to poor performance in highly structured probabilistic programs with complex computational flows (Ambrogioni et al., 2021). For example, consider an inference problem defined by a highly parameterized natural language model (Brown et al., 2020). A typical problem of this kind is controlled language generation, where the goal is to steer the model to general text conditioned on its past output and future targets (e.g. the presence of some specific words). The controlled network can be formulated in probabilistic programming terms as a posterior distribution that is parameterized by the variational program. In a problem like this, most of the complexity of the posterior is encoded in the structure of the prior probabilistic program (the language model). The model just needs to be steered in the right direction. However, normalizing flow completely disregards this structure and has to relearn all the structure embedded in the prior.

This problem was discussed in a recent work that introduced a form of automatic structured variational inference (ASVI) with an automated variational program that incorporates the prior probabilistic program as a special case. ASVI has been shown to have very good performance in time series analysis and deep learning problems in spite of being very parsimoniously parameterized compared to normalizing flows (Ambrogioni et al., 2021). However, the convex-update families used in ASVI have two main limitations: I) they cannot model statistical dependencies arising from colliders in the graphical model (the "explaining away" effect) and II) they constrain the conditional posterior distributions to be in the same family as the prior. In this paper we integrate ASVI and normalizing flows to obtain a new form of automatic VI, which we refer to as *cascading flows* (CF). The approach incorporates the forward-pass of the probabilistic program while being capable of modeling collider dependencies and arbitrarily complex conditional distributions. To this end, we make use of a novel highway flow architecture. CF respects the following design principles:

- **Automation:** It can be constructed using a fully automatic algorithm that takes the prior probabilistic program as input.
- **Locality:** It is constructed by locally transforming each conditional distribution in the conditional prior so that it inherits the graphical structure of the input program (with potentially some extra coupling). This ensures scalability and modularity.
- **Prior information flow:** It embeds the forward-pass of the probabilistic program.

2. Related work

Automatic algorithms for VI date back to VIBES (Bishop et al., 2002; Bishop and Winn, 2003). These older works are based on self-consistency equations and variational message passing (Winn and Bishop, 2005). The use of gradient-based automatic VI for probabilistic programming was pioneered in (Wingate and Weber, 2013) using a mean field approach. Automatic differentiation VI (ADVI) perfected this approach by exploiting the automatic-differentiation tools developed in the deep learning community (Kucukelbir et al., 2017). ADVI uses fixed invertible transformations to map each variable into an unbounded domain where the density can be modeled with either univariate or multivariate normal distributions. Structured stochastic VI (Hoffman and Blei, 2015) exploits the local conjugate update rules of locally conjugate models. This has the major benefit of preserving the local structure of the probabilistic program but it severely limits the class of possible input programs since local conjugacy is a stringent condition. Automatic structured VI (ASVI) lifts this constraint as it applies a

trainable convex-update rule to non-conjugate links (Ambrogioni et al., 2021). Several other papers introduced new techniques for constructing structured variational programs. Copula VI introduces the use of vine copulas to model the statistical coupling between the latent variables in the program (Tran et al., 2015). Similarly, hierarchical VI uses auxiliary latent variables to achieve the same goal (Ranganath et al., 2016). Neither of those approaches are truly automatic however, as they do not prescribe a unique way to construct the coupling functions.

Normalizing flows are very expressive and highly parameterized models that use deep learning components to approximate complex multivariate densities (Rezende and Mohamed, 2015; Dinh et al., 2014; 2017). Normalizing flows can approximate arbitrarily complex densities and can therefore be turned into powerful automatic VI methods (Rezende and Mohamed, 2015; Kingma et al., 2016). However, conventional normalizing flows do not exploit the input probabilistic program. Structured conditional continuous normalizing flows are a newly introduced class of flows constrained to have the same conditional independence structure of the true posterior (Weilbach et al., 2020). However, these flow architectures only inherit the graphical structure of the input program, without incorporating the forward pass (i.e. the form of the conditional distributions and link functions).

Modern deep probabilistic programming frameworks provide automatic implementation of many of variational programs. For example, PyMC (Patil et al., 2010) provides an implementation of ADVI, Pyro (Bingham et al., 2019) implements the automatic construction of mean field, multivariate normal and normalizing flow programs (i.e. variational guides) and TensorFlow probability also offers an automatic implementation of ASVI (Dillon et al., 2017; Ambrogioni et al., 2021).

Our new highway flow architecture is a key component of our new approach and it is inspired by highway networks (Srivastava et al., 2015). To the best of our knowledge, highway flows are the first form of residual-like flows with a tractable Jacobian determinant that can be expressed in closed form. Existing residual flows express the log determinant of the Jacobian as an infinite series that needs to be estimated using a "Russian roulette" estimator (Chen et al., 2019). We achieve tractability by applying a highway operator to each single layer of the network.

The use of auxiliary variables has proved to be an effective way to improve the expressiveness of Normalizing Flows by operating in an augmented space (Huang et al., 2020; Cornish et al., 2020; Caterini et al., 2020). Our use of auxiliary variables is similar to the approach adopted in (Dupont et al., 2019) and (Weilbach et al., 2020). The bound we use for training the augmented model was derived in (Ranganath

et al., 2016).

Structured variational programs have the most clear benefits in timeseries analysis problems (Eddy, 1996; Foti et al., 2014; Johnson and Willsky, 2014; Karl et al., 2016; Fortunato et al., 2017). Our approach, together with ASVI, differs from the conventional timeseries approach by exploiting the structure of both temporal and non-temporal variables in the model.

3. Preliminaries

In this section we introduce the notation used in the rest of the paper and discuss in detail the methods that form the basis for our work.

3.1. Differentiable probabilistic programs

We denote the values of (potentially multivariate) random variables using lowercase notation and arrays of random variables using boldface notation. We consider probabilistic programs constructed as a chain of conditional probabilities over an array of random variables $\mathbf{x} = (x_1, \dots, x_N)$. We denote the array of values of the parents of the j -th variable as $\boldsymbol{\pi}_j \subseteq \{x_i\}_{i \neq j}$, which is a sub-array of parent variables such that the resulting graphical model is a directed acyclic graph. Using this notation, we can express the joint density of a probabilistic program as follows:

$$p(\mathbf{x}) = \prod_{j=1}^N \rho_j(x_j | \theta_j(\boldsymbol{\pi}_j)), \quad (1)$$

where $\theta_j(\boldsymbol{\pi}_j)$ is a link function that maps the values of the parents of the j -th variable to the parameters of its density $\rho_j(\cdot | \cdot)$. For example, if the density is Gaussian the link function outputs the value of its mean and scale given the values of the parents. Note that θ_j is a constant when the array of parents is empty.

3.2. Convex-update variational programs

An automatic structured variational inference method provides an algorithm that takes a probabilistic program as input and outputs a variational program $q(\mathbf{x})$. Convex-update variational programs (Ambrogioni et al., 2021) are defined by the following transformation:

$$p(\mathbf{x}) \xrightarrow{\text{CU}} q(\mathbf{x}) = \prod_j \mathcal{U}_{\lambda_j}^{\alpha_j} \rho_j(x_j | \theta_j(\boldsymbol{\pi}_j)), \quad (2)$$

with convex-update operator

$$\mathcal{U}_{\lambda_j}^{\alpha_j} \rho_j(x_j | \theta_j(\boldsymbol{\pi}_j)) = \rho_j(x_j | \lambda_j \odot \theta_j(\boldsymbol{\pi}_j) + (1 - \lambda_j) \odot \alpha_j), \quad (3)$$

where λ_j and α_j are (potentially vector-valued) learnable parameters and \odot is the element-wise product. This reduces to the prior probabilistic program for $\lambda_j \rightarrow 0$ and to a mean field variational program for $\lambda_j \rightarrow 1$. Convex-update ASVI has the advantage of preserving the forward-pass of the probabilistic program, which is often a major source of statistical coupling in the posterior distribution. The variation program can be trained to fit the posterior by minimizing the evidence lower bound (ELBO) by stochastic gradient descent (SGD) on the differentiable parameters λ_j and α_j .

3.3. Variational inference with normalizing flows

Normalizing flows express the density of arbitrary multivariate distributions as the push forward of a known base distribution through a differentiable invertible function (i.e. a diffeomorphism). The base distribution is usually a spherical normal distribution, hence the "normalizing" in the name. Consider a d -variate probability density $p_0(z)$ and a diffeomorphism $\Psi : \mathbb{R}^d \rightarrow \mathbb{R}^d$. We can derive the density of the new random variable $x = \Psi(z)$ using the change of variable formula

$$p_X(x) = |\det J(\Psi^{-1}(x))| p_0(\Psi^{-1}(x)), \quad (4)$$

where Ψ^{-1} is the inverse transformation and $J(z)$ is the Jacobi matrix of Ψ . We can write this more compactly using the push-forward operator \mathcal{T}_Ψ associated to Ψ that maps the density $p_0(z)$ to the density $p_X(x)$ given in Eq. 4: $p_X(x) = \mathcal{T}_\Psi[p_0(\cdot)](x)$.

Now consider a flexible family of functions $\Psi^w(x)$ parameterized by the "weights" w . We can use Eq. 13 to approximate arbitrarily complicated target densities by training w by SGD. Given a probabilistic program $p(x, y)$ with y observed, we can approximate the posterior $p(x | y)$ by minimizing the evidence lower bound:

$$\text{ELBO}(w) = \mathbb{E}_w \left[\log \frac{p(x, y)}{|\det J_w(\Psi^{-1}(x))| p_0(\Psi^{-1}(x))} \right] \quad (5)$$

where the expectation is taken with respect to the transformed distribution $\mathcal{T}_{\Psi^w}[p_0(\cdot)](x)$. Since Ψ is differentiable, we can obtain the reparameterization gradient for the ELBO as follows:

$$\nabla_w \text{ELBO}(w) = \mathbb{E}_0 \left[\nabla_w \log \frac{p(\Psi^w(z), y)}{|\det J_w(z)| p_0(z)} \right] \quad (6)$$

where the expectation is now taken with respect to the fixed distribution $p_0(z)$. In the following we refer to this approach as *global flow* (GF) since the transformation is applied synchronously to all variables in the model.

4. Variational inference with cascading flows

Now that we have covered the relevant background, we will introduce our new approach that combines ASVI with

normalizing flows.

4.1. Cascading flows variational programs

The convex-update parameterization poses strong constraints on the form of the conditional posteriors. On the other hand, normalizing flow programs do not embed the potentially very complex structure of the input program. In this paper we combine the two approaches by replacing the convex-update operator in Eq. 2 with the push-forward operator used in normalizing flows. More precisely, consider a family of diffeomorphisms $\Psi^w = (\Psi_1^{w_1}, \dots, \Psi_N^{w_N})$ depending on the array of parameters $w = (w_1, \dots, w_N)$. Using these transformations, we can define the *cascading flows* variational program associated with a program defined by Eq. 1:

$$p(x) \mapsto_{\text{CF}} q_w(x) = \prod_j^N \mathcal{T}_j^w [\rho_j(\cdot | \theta_j(\pi_j))](x_j), \quad (7)$$

where we denoted the push-forward operator induced by the diffeomorphism $\Psi_j^{w_j}$ (as defined by Eq. 4) as \mathcal{T}_j^w to simplify the notation. A cascading flows variational program can be trained by SGD using gradient estimators of the ELBO with the reparameterization given in Eq. 6.

The name cascading flows comes from the fact that, since the transformation is performed locally for each conditional distribution, its effects cascade down the variational model through the (transformed) conditional dependencies of the prior program. However, in order to preserve the forward-pass of the probabilistic program it is crucial to use a parameterized diffeomorphism that can be initialized around the identity function and whose deviation from the identity can be easily controlled. Specifically, we opt for transformations of the following form:

$$\Psi_j^{w_j}(x) = \gamma x + (1 - \gamma) f_j(x; w_j), \quad (8)$$

where $\gamma \in (0, 1)$. The resulting architecture of a cascading flows variational program is visualized in Fig 1. From the diagram it is clear that a cascading flow is constructed by inserting transformation layers within the architecture of the probabilistic program. Transformations of the form given in Eq. 8 with tractable Jacobians are not currently present in the normalizing flow literature. Therefore in the next subsection we will introduce a new flow architecture of this form, which we named *highway flow*.

4.2. Highway flow architecture

Here we introduce highway flow blocks gated by a learnable scalar gate parameter. As also visualized in Fig. 2, a highway flow block is comprised by three separate highway layers:




Figure 1. Diagram of a cascading flows architecture. A) Detail of architecture of an input probabilistic program. B) Detail of the associated cascading flows architecture.




Figure 2. Diagram of a highway flow block.

1. Upper triangular highway layer:

$$l_U(z; U, \lambda) = \lambda z + (1 - \lambda) (Uz + b_U) \quad (9)$$

$$\log \det J_U = \sum_k \log (\lambda + (1 - \lambda) U_{kk}) \quad (10)$$

2. Lower triangular layer:

$$l_L(z; L, \lambda) = \lambda z + (1 - \lambda) (Lz + b_L) \quad (11)$$

$$\log \det J_L = \sum_k \log (\lambda + (1 - \lambda) L_{kk}) \quad (12)$$

3. Highway activation functions:

$$f(z; \lambda) = \lambda z + (1 - \lambda) g(z) \quad (13)$$

$$\log \det \frac{df(x_k)}{dx} = \sum_k \log \left(\lambda + (1 - \lambda) \frac{dg(x_k)}{dx} \right) \quad (14)$$

where U is an upper-triangular matrix with positive-valued diagonal, L is a lower-triangular matrix with ones in the diagonal and $g(x)$ is a differentiable non-decreasing activation function. A highway flow layer is a composition of these three types of layers:

$$l_h(\cdot) = f \circ l_L \circ l_U(\cdot). \quad (15)$$

A highway network is a sequence of M highway layers with a common gate parameter and different weights and biases. Note that a highway flow can be expressed in the form given in Eq. 8 by defining γ as follows:

$$\gamma = \lambda^{3M}, \quad (16)$$

which is clearly equal to one (zero) when λ is equal to one (zero).

4.3. Auxiliary variables and infinite mixtures of flows

The expressiveness of a normalizing flow is limited by the invertible nature of the transformation. This is particularly problematic for low-dimensional variables since the expressivity of many flow architectures depends on the dimensionality of the input (Papamakarios et al., 2019). This limitation can be a major shortcoming in cascading flows as they are particularly useful in highly structured problems characterized by a large number of coupled low-dimensional variables. Fortunately, we can remedy this limitation by introducing a set of auxiliary variables. For each variable x_j , we create a D -dimensional variable ϵ_j following a base distribution $p_j(\epsilon_j)$. We can now use an augmented diffeomorphism $\hat{\Psi}_j^{w_j}(x_j, \epsilon_j)$ and define the joint posterior over both x_j and ϵ_j :

$$q(x_j, \epsilon_j | \pi_j) = \hat{T}_j^w [\rho_j(\cdot | \theta_j(\pi_j)) p_j(\cdot)](x_j, \epsilon_j), \quad (17)$$

where the push-forward operator now transforms the (independent) joint density of x_j and ϵ_j . The conditional variational posterior is then obtained by marginalizing out ϵ_j :

$$q(x_j | \pi_j) = \int q(x_j, \epsilon_j | \pi_j) d\epsilon_j. \quad (18)$$

This infinite mixture of flows is much more capable of modeling complex and potentially multi-modal distributions.

The use of latent mixture models for VI was originally introduced in (Ranganath et al., 2016). Given a set of observations \mathbf{y} , the ELBO of any mixture variational probabilistic program can be lower bounded using Jensen's inequality:

$$\mathbb{E}_{\mathbf{x}} \left[\log \frac{p(\mathbf{x}, \mathbf{y})}{\int q(\mathbf{x}, \epsilon) d\epsilon} \right] \geq \underbrace{\mathbb{E}_{\mathbf{x}, \epsilon} \left[\log \frac{p(\mathbf{x}, \mathbf{y}) r(\epsilon)}{q(\mathbf{x}, \epsilon)} \right]}_{\text{Augmented ELBO}}, \quad (19)$$

where $q(\epsilon) = \int q(\mathbf{x}, \epsilon) d\mathbf{x}$ is the marginal variational posterior and $r(\epsilon)$ is an arbitrarily chosen distribution over the auxiliary variables.

This result justifies the use of the augmented ELBO for training the variational posterior. In practice, the distribution $r(\epsilon)$ can be chosen to be a mean field variational density. The gap between the mixture ELBO (Eq. 19) and the augmented

ELBO can be reduced by training these mean field parameters by minimizing the augmented ELBO such as to match $q(\epsilon)$ (Ranganath et al., 2016).

The highway flow architecture can be easily adapted to the augmented variable space. Since there is no need to gate the flow of the auxiliary variables, the scalar λ in Eq. 9, Eq. 11 and Eq. 13 should be replaced by the vector \mathbf{l} whose entries corresponding to the original variables are equal to λ while the entries corresponding to the auxiliary variables are equal to zero.

4.4. Modeling collider dependencies with auxiliary coupling




Figure 3. Backward auxiliary coupling. A) Example graphical model. B) Cascading flows model coupled with reversed local linear Gaussian model.

So far, we outlined a model that has the same conditional independence structure of the prior and is therefore incapable of modeling dependencies arising from colliders. Fortunately, the local auxiliary variables that we introduced in order to increase the flexibility of the local flows can assume a second role as "input ports" that can induce non-local couplings. In fact, by coupling their auxiliary variables we can induce statistical dependencies between any subset of variables in the model. Importantly, this coupling can take a very simple form since additional complexity can be modeled by the highway flows. We will now introduce a local auxiliary model with a non-trivial conditional independence structure inspired by the forward-backward algorithm. The graphical structure of the posterior probabilistic program is a subgraph of the moralization of the prior graph (Bishop, 2006). Therefore, in order to be able to capture all the dependencies in the posterior, it suffices to couple the auxiliary variables by reversing the arrows of the original prior model. In fact, all the parents of every node get coupled by the reversed arrows. As an example, consider the following probabilistic model (Fig. 3A): $\rho_1(x_1 | x_2, x_3)\rho_2(x_2)\rho_3(x_3 | x_4)\rho_4(x_4)$. We couple the auxiliary variables by reversing the arrows, as

Table 1. Predictive and latent log-likelihood (forward KL) of variational timeseries models. Error are SEM estimated over 10 repetitions.

		CF	ASVI	MF	GF	MVN	CF (non-res)
BR-r	Pred	-2.27 ± 0.26	-2.23 ± 0.21	-3.79 ± 0.82	-2.81 ± 0.56	-2.88 ± 0.53	-3.33 ± 0.65
	Latent	-1.48 ± 0.19	-1.45 ± 0.14	-4.02 ± 0.63	-2.41 ± 0.52	-2.02 ± 0.48	-3.63 ± 0.74
BR-c	Pred	1.61 ± 0.18	1.45 ± 0.14	1.04 ± 0.03	2.00 ± 0.29	1.02 ± 0.03	1.31 ± 0.18
	Latent	-1.53 ± 0.21	-1.55 ± 0.19	-5.78 ± 0.89	-2.06 ± 0.53	-2.82 ± 0.77	-5.07 ± 0.85
LZ-r	Pred	-2.89 ± 0.17	-4.48 ± 0.60	-8.26 ± 0.28	-8.03 ± 0.37	-8.24 ± 0.29	-8.25 ± 0.27
	Latent	-2.39 ± 0.45	-4.38 ± 0.67	-10.28 ± 0.18	-9.44 ± 0.20	-9.45 ± 0.22	-10.00 ± 0.18
LZ-c	Pred	5.10 ± 0.52	0.92 ± 0.03	0.90 ± 0.003	0.86 ± 0.15	0.89 ± 0.001	0.88 ± 0.04
	Latent	-4.19 ± 0.66	-7.47 ± 0.30	-9.89 ± 0.19	-8.71 ± 0.32	-8.58 ± 0.34	-9.59 ± 0.29
PD-r	Pred	-3.19 ± 0.22	-3.25 ± 0.11	-4.42 ± 0.22	-3.84 ± 0.28	-4.30 ± 0.22	-4.29 ± 0.25
	Latent	-2.32 ± 0.19	-3.14 ± 0.12	-9.12 ± 0.29	-4.16 ± 0.33	-7.72 ± 0.30	-8.27 ± 0.36
PD-c	Pred	1.97 ± 0.07	1.65 ± 0.06	0.86 ± 0.003	01.07 ± 0.02	1.09 ± 0.02	0.96 ± 0.01
	Latent	-2.77 ± 0.18	-3.09 ± 0.15	-8.40 ± 0.43	-6.20 ± 0.40	-7.45 ± 0.42	-8.41 ± 0.43
RNN-r	Pred	-1.68 ± 0.05	-2.30 ± 0.18	-5.20 ± 0.94	-1.60 ± 0.09	-4.47 ± 0.92	-1.97 ± 0.21
	Latent	-1.34 ± 0.33	-1.95 ± 0.35	-10.30 ± 0.20	-6.39 ± 1.27	-6.61 ± 0.50	-10.47 ± 0.22
RNN-c	Pred	5.77 ± 1.40	1.05 ± 0.06	0.81 ± 0.03	2.81 ± 0.36	0.86 ± 0.02	1.39 ± 0.04
	Latent	-2.30 ± 0.61	-2.05 ± 0.32	-10.22 ± 0.29	-10.75 ± 0.15	-10.22 ± 0.29	-11.22 ± 0.04

formalized in the graphical model $p_1(\epsilon_1)p_2(\epsilon_2, | \epsilon_1)p_3(\epsilon_3, | \epsilon_1)p_4(\epsilon_4 | \epsilon_3)$. The graphical structure of the resulting augmented model is shown in Fig. 3B. The auxiliary variables can be coupled in a very simple manner since the diffeomorphisms can add arbitrary complexity to the final joint density. Here we use a simple linear normal model:

$$\epsilon_k | \mathbf{v}_k = \sum_{j=1}^K a_j \odot v_j + a_0 \odot \xi_k, \quad (20)$$

where $a_j \geq 0$ and $\sum_n a_n = 1$. In this expression, ξ_k is a standard Gaussian vector and $\mathbf{v}_k = (v_1, \dots, v_K)$ is an array containing all the parents auxiliary variables (the auxiliaries of the children of x_k in the original graph).

4.5. Inference amortization

We will now discuss a strategy to implement automatic inference amortization in cascading flows models. The problem of inference amortization is closely connected to the problem of modeling collider dependencies as the latter arise from the backward messages originating from observed nodes (Bishop, 2006). It is therefore not a coincidence that we can implement both collider dependencies and amortization by shifting the statistics of the local auxiliary variables. This results in a very parsimonious amortization scheme where data streams are fed locally to the various nodes and propagate backward through the linear network of auxiliary Gaussian variables. We denote the set of auxiliary variables associated to the children of x_k in the input probabilistic program as \mathbf{v}_k . Furthermore, we denote the observed value of x_k as y_k and the set of observed values of the j -th child of x_k as $\mathbf{y}^k = (y_1^k, \dots, y_K^k)$. The

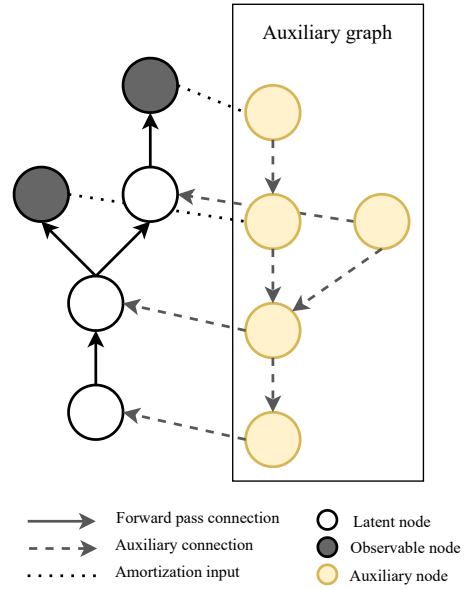


Figure 4. Visualization of the backward amortization algorithm.

amortized auxiliary model can then be obtained as a modification of Eq. 20: $p_0(\epsilon_k | \mathbf{v}_k) = \mathcal{N}(m_k, \sigma^2 I)$ with $\epsilon_k | \mathbf{v}_k = \mathcal{B}^{(k)}[y_k] + \sum_{j=1}^K a_j \odot v_j + a_0 \odot \xi_k$, where $\mathcal{B}^{(k)}$ are learnable affine transformations (i.e. linear layers in deep learning lingo).

5. Automation

Cascading flow variational programs can be constructed automatically given an input probabilistic program. We

provide an open-source implementation of this algorithm in TensorFlow Probability (Dillon et al., 2017). The code for the spacial case without amortization and backward auxiliary coupling is shown in Figure 5.

```
def build_cascading_flows_surrogate_posterior(prior):
    q_vars = {}
    @tfed.JointDistributionCoroutineAutoBatched
    def cascading_flows_surrogate_posterior():
        # Step the model to yield the first RV.
        prior_gen = prior._coroutine_fn()
        rv = next(prior_gen)
        while True: # Run model to termination.
            # If this is a new RV, initialize variables.
            if rv.name not in q_vars:
                q_vars[rv.name] =
                    build_cascading_flows_bijector(
                        width=rv.width,
                        num_highway_flows_blocks=3)
            # Sample from the Cascading Flows.
            q_sample = yield tfed.TransformedDistribution(
                distribution=rv,
                bijector=q_vars[rv.name])
            # Step to the next RV.
            rv = prior_gen.send(q_sample)
    return cascading_flows_surrogate_posterior
```

Figure 5. TFP code for automatically constructing a Cascading Flows variational program given the prior program.

6. Experiments

Our experiments are divided into three sections. In the first section we focus on highly structured timeseries problems exhibiting linear and non-linear dynamics. Bayesian problems of this nature have relevance for several scientific fields such as physics, chemistry, biology, neuroscience and population dynamics. In the second section we compare the capacity of several variational programs to model dependencies arising from colliders in the graphical structure. In fact, the inability of modeling collider dependencies was the main limitation of the original convex-update ASVI family (Ambrogioni et al., 2021). Finally, in the third section we test the performance of the automatic amortization scheme. All the code used in the experiments is included in the supplementary material and documented in Supplementary A. **Architectures:** In all experiments, the CF architectures were comprised of three highway flow blocks with softplus activation functions in each block except for the last which had linear activations. CF programs use an independent network of this form for each latent variable in the program. Each variable was supplemented with 10 auxiliary variables, the width of each network was therefore equal to the dimensionality of the variable plus 10. Weights and biases were initialized from centered normal distributions with scale 0.01. The λ variable was defined independently for each network as the logistic sigmoid of a learnable parameter l ,

which was initialized as 4 in order to keep the variational program close to the input program. For each variable, the posterior auxiliary distributions $r(\epsilon)$ (see Eq. 19) were spherical normals parameterized by mean and standard deviation (i.e. Gaussian mean field).




Figure 6. Example of qualitative results in the Lorenz system (first coordinate). The gray shaded area is observed through noisy measurements. The dashed lines are the ground-truth timeseries while the thin colored lines are samples from the trained variational program.

6.1. Timeseries analysis

We consider discretized SDEs, defined by conditional densities of the form:

$$\rho_t(x_{t+1} | x_t) = \mathcal{N}(x_{t+1}; \mu(x_t, t)dt, \sigma^2(x_t, t)dt) \quad (21)$$

where $\mu(x_t, t)$ is the drift function and $\sigma(x_t, t)$ is the volatility function. Specifically, we used Brownian motions (BR: $\mu(x) = x, \sigma^2(x) = 1$), Lorenz dynamical systems (LZ: $\mu(x_1, x_2, x_3) = (10(x_2 - x_1), x_1(28 - x_3) - x_2, x_1x_2 - 8/3x_3), \sigma^2 = 2$), Lotka-Volterra population dynamics models (PD: $\mu(x_1, x_2) = (0.2x_1 - 0.02x_1x_2, 0.1x_1x_2 - 0.1x_2), \sigma^2 = 3$) and recurrent neural dynamics with randomized weights (RNN, see Supplementary B). The initial distributions were spherical Gaussians (see Supplementary B for the parameters). For each dynamical model we generated noisy observations using Gaussian emission models (denoted with an "r" in the tables):

$$p(y_t | x_t) = \mathcal{N}([x_t]_1, \sigma_{lk}^2) \quad (22)$$

or Bernoulli logistic models ("c" in the tables):

$$p(y_t | x_t) = \text{Bernoulli}(g(k[x_t]_1)) \quad (23)$$

where $[x_t]_1$ denotes the first entry of the vector x_t , σ_{lk} is the standard deviation of the Gaussian emission model,

Table 2. Multivariate latent log-likelihood (forward KL) of variational binary tree models. Error are SEM estimated over 15 repetitions.

	CF	ASVI	MF	GF	MN
Linear-2	1.24 ± 0.01	0.49 ± 0.05	0.6 ± 0.01	0.97 ± 0.01	1.00 ± 0.01
Linear-4	5.42 ± 0.35	3.56 ± 0.26	1.68 ± 0.02	6.56 ± 0.04	6.52 ± 0.01
Tanh-2	1.20 ± 0.01	-7.86 ± 3.47	0.24 ± 0.01	0.90 ± 0.01	1.15 ± 0.01
Tanh-4	3.29 ± 3.17	-55.44 ± 4.53	-43.24 ± 8.87	10.32 ± 0.05	5.37 ± 2.61

$g(\cdot)$ is the logistic sigmoid function and k is a gain factor. All the numeric values of the parameters in different experiments are given in Supplementary B. For each task we evaluate the performance of the cascading flows programs (CF) against a suite of baseline variational posteriors, including convex-update ASVI (Ambrogioni et al., 2021), mean field (MF) and multivariate normal (MN) (Kucukelbir et al., 2017). We also compare our new cascading flows family with the more conventional global normalizing flow (GF) approach where a single diffeomorphism transforms all the latent variables in the probabilistic program into a spherical normal distribution (Rezende and Mohamed, 2015). For the sake of comparison, we adopt an architecture identical to our highway flow model (including 10 auxiliary variables) except for the absence of the highway gates. Note however that the global architecture has much greater width as it is applied to all variables in the program. Finally, we compare the performance of CF with a modified version of CF that does not use highway gates. Ground-truth multivariate timeseries $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_T)$ were sampled from the generative model together with simulated first-half observations $\mathbf{y}_{1:T/2} = (y_1, \dots, y_{T/2})$ and second-half observations $\mathbf{y}_{T/2:T} = (y_{T/2+1}, \dots, y_T)$. All the variational models were trained conditioned only on the first half observations. Performance was assessed using two metrics. The first metric is the average marginal log-probability of the ground-truth given the variational posterior: $\frac{1}{TJ} \sum_t \sum_j \log q([\tilde{x}_t]_j | \mathbf{y}_{1:T/2})$, where J is the dimensionality of the ground-truth vector \tilde{x}_t . The marginal posterior density $q([\tilde{x}_t]_j | \mathbf{y}_{1:T/2})$ was estimated from 5000 samples drawn from the variational program using Gaussian kernel density estimation with bandwidth $0.9\hat{\sigma}N^{-1/5}$, where $\hat{\sigma}$ is the empirical standard deviation of the samples. Our second metric is $\log p(\mathbf{y}_{T/2:T} | \mathbf{y}_{1:T/2})$: the predictive log-probability of the ground-truth observations in the second half of the timeseries given the observations in the first half, estimated using kernel density estimation from 5000 predicted observation samples drawn from the exact likelihood conditioned on the samples from the variational program. Each experiment was repeated 10 times. In each repetition, all the variational programs were re-trained for 8000 iterations (enough to ensure convergence in all methods) given new sampled ground-truth timeseries and observations. The reported results are the mean and SEM of the two metrics over these 10 repetitions. **Results:** The

results are given in Table 1. The cascading flows family (CF) outperforms all other approaches by a large margin in all non-linear tasks. Convex-update ASVI consistently has the second highest performance, easily outperforming the highly parameterized global normalizing flow (GF). This result is consistent with (Ambrogioni et al., 2021), where ASVI was shown to outperform inverse autoregressive flow architectures (Kingma et al., 2016). Finally, the CF model without highway gates (CF non-res) has very low performance, proving the crucial importance of our new highway flow architecture. A qualitative comparison (LZ-r and LZ-c) between CF and ASVI is shown in Fig. 6.

6.2. Binary tree experiments

In our second experiment we test the capability of the variational programs to model statistical dependencies arising from colliders. We tested this in a Gaussian binary tree model where the mean of a scalar variable x_j^d in the d -th layer is a function of two variables in the $d-1$ -th layer: $x_j^d \sim \mathcal{N}(\text{link}(\pi_{j1}^{d-1}, \pi_{j2}^{d-1}), \sigma^2)$, where π_{j1}^{d-1} and π_{j2}^{d-1} are the two parents of x_j^d . All the 0-th layer variables have zero mean. The inference problem consists of observing the only variable in the last layer and computing the posterior of all the other variables. We considered a setting with linear coupling $\text{link}(x, y) = x - y$ and tanh coupling $\text{link}(x, y) = \tanh(x) - \tanh(y)$ with 2 and 4 layers. Performance was estimated using the posterior log-probability of the ground-truth sample estimated using a multivariate Gaussian approximation that matches the empirical mean and covariance matrix of 5000 samples drawn from the trained variational program. We compared performance with ASVI, mean-field (MF), global flow (GF) and multivariate normal (MN). Note that neither ASVI nor MF can model collider dependencies. Full details are given in Supplementary C. The reported results are the mean and SEM of the two metrics over 15 repetitions. **Results:** The results are given in Table 2. The CF model always outperforms ASVI and MF, proving that it can capture collider dependencies. CF has the highest performances in shallow models but lower performance than GF and MN in deeper models.

6.3. Inference amortization

Finally, we evaluate the performance of the amortization scheme in non-linear timeseries inference problems. We

Table 3. Latent log-likelihood (forward KL) of amortized time-series models. Error are SEM estimated over 50 repetitions.

	CF	MF	GF
PD	-2.8 ± 0.2	-8.0 ± 0.2	-5.1 ± 0.2
LZ	-2.7 ± 0.2	-10.0 ± 0.2	-9.1 ± 0.2

used the population dynamics (PD) and the Lorentz (LZ) systems. The CF program was amortized using the approach described in Sec. 4.5. The MF and normalizing flow baselines were amortized using inference networks (see Supplementary D). All the details were identical to the timeseries experiments except for the fact that all time points were observed. Performance was again quantified as the average marginal log-probability of the ground-truth kernel density estimated from the samples from the amortized variational programs. Each amortized program was trained once and tested on 50 new sampled timeseries/observations pairs. The reported results are the mean and SEM of these 50 repetitions. **Results:** The results are reported in Table 3. Like in the non-amortized case, CF greatly outperforms both MF and GF in both timeseries inference problems.

7. Discussion

Our CF approach solves the two main limitations of ASVI as it can model collider dependencies and can be amortized automatically. We show that cascading flows programs have extremely high performance in structured inference problems. However, their ability to capture deep collider dependencies is lower than normalizing flows. In our opinion, designing principled ways to model collider dependencies is a promising research direction.

References

- B. Milch, B. Marthi, S. Russell, D. Sontag, D. L. Ong, and A. Kolobov. Blog: Probabilistic models with unknown objects. *Statistical Relational Learning*, page 373, 2007.
- T. Sato. Prism: A symbolic-statistical modeling language. In *Proceeding of the Intl. Joint Conference on Artificial Intelligence*, 1997.
- K. Kersting and L. De Raedt. Bayesian logic programming: Theory and tool. *Statistical Relational Learning*, page 291, 2007.
- A. Pfeffer. Ibal: A probabilistic rational programming language. In *Proceeding of the Intl. Joint Conference on Artificial Intelligence*, 2001.
- S. Park, F. Pfenning, and S. Thrun. A probabilistic language based upon sampling functions. *ACM SIGPLAN Notices*, 40(1):171–182, 2005.
- N. Goodman, V. Mansinghka, D. M. Roy, K. Bonawitz, and J. B. Tenenbaum. Church: A language for generative models. *arXiv preprint arXiv:1206.3255*, 2012.
- D. Wingate, A. Stuhlmüller, and N. Goodman. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2011.
- A. Patil, D. Huard, and C. J. Fonnesebeck. Pymc: Bayesian stochastic modelling in python. *Journal of Statistical Software*, 35(4):1, 2010.
- J. V Dillon, I. Langmore, D. Tran, E. Brevdo, S. Vasudevan, D. Moore, B. Patton, A. Alemi, M. Hoffman, and R. A. Saurous. Tensorflow distributions. *arXiv preprint arXiv:1711.10604*, 2017.
- E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1):973–978, 2019.
- D. Tran, M. D. Hoffman, R. A. Saurous, E. Brevdo, K. Murphy, and D. M. Blei. Deep probabilistic programming. *arXiv preprint arXiv:1701.03757*, 2017.
- D. Tran, A. Kucukelbir, A. B. Dieng, M. Rudolph, D. Liang, and D. M. Blei. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787*, 2016.
- C. Bishop, D. Spiegelhalter, and J. Winn. Vibes: A variational inference engine for bayesian networks. *Neural Information Processing Systems*, 15:793–800, 2002.
- A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. Automatic differentiation variational inference. *The Journal of Machine Learning Research*, 18(1):430–474, 2017.
- C. Bishop and J. Winn. Structured variational distributions in vibes. In *International Conference on Artificial Intelligence and Statistics*, 2003.
- D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.
- G. Papamakarios, E. Nalisnick, Danilo J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- L. Ambrogioni, K. Lin, E. Fetig, S. Vikram, M. Hinne, D. Moore, and M. van Gerven. Automatic structured

- variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Neural Information Processing Systems*, 2020.
- J. Winn and C. M. Bishop. Variational message passing. *Journal of Machine Learning Research*, 6(Apr):661–694, 2005.
- D. Wingate and T. Weber. Automated variational inference in probabilistic programming. *arXiv preprint arXiv:1301.1299*, 2013.
- M. D. Hoffman and D. M. Blei. Structured stochastic variational inference. In *Proceedings of the Artificial Intelligence and Statistics*, pages 361–369, 2015.
- D. Tran, D. Blei, and E. M. Airoldi. Copula variational inference. In *Neural Information Processing Systems*, 2015.
- R Ranganath, D Tran, and D Blei. Hierarchical variational models. In *International Conference on Machine Learning*, pages 324–333, 2016.
- L. Dinh, D. Krueger, and Y. Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real NVP. In *International Conference on Learning Representations*, 2017.
- D. P Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Neural Information Processing Systems*, 2016.
- C Weilbach, B Beronov, F Wood, and W Harvey. Structured conditional continuous normalizing flows for efficient amortized inference in graphical models. In *International Conference on Artificial Intelligence and Statistics*, pages 4441–4451. PMLR, 2020.
- R K Srivastava, K Greff, and J Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- R. T.Q. Chen, J. Behrmann, D. K. Duvenaud, and J. Jacobsen. Residual flows for invertible generative modeling. In *Neural Information Processing Systems*, pages 9916–9926, 2019.
- C. Huang, L. Dinh, and A. Courville. Augmented normalizing flows: Bridging the gap between generative flows and latent variable models. *arXiv preprint arXiv:2002.07101*, 2020.
- R. Cornish, A. Caterini, G. Deligiannidis, and A. Doucet. Relaxing bijectivity constraints with continuously indexed normalising flows. In *International Conference on Machine Learning*. PMLR, 2020.
- A. Caterini, R. Cornish, D. Sejdinovic, and A. Doucet. Variational inference with continuously-indexed normalizing flows. *arXiv preprint arXiv:2007.05426*, 2020.
- E Dupont, A Doucet, and Y W Teh. Augmented neural ODEs. In *Neural Information Processing Systems*, pages 3140–3150, 2019.
- S. R. Eddy. Hidden Markov models. *Current Opinion in Structural Biology*, 6(3):361–365, 1996.
- N. Foti, J. Xu, D. Laird, and E. Fox. Stochastic variational inference for hidden Markov models. In *Neural Information Processing Systems*, 2014.
- M. Johnson and A. Willsky. Stochastic variational inference for Bayesian time series models. In *International Conference on Machine Learning*, 2014.
- M. Karl, M. Soelch, Ju. Bayer, and P. Van der Smagt. Deep variational Bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- M. Fortunato, C. Blundell, and O. Vinyals. Bayesian recurrent neural networks. *arXiv preprint arXiv:1704.02798*, 2017.
- C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.