
Automatic structured variational inference

Luca Ambrogioni¹

Max Hinne¹

Kate Lin²

Dave Moore²

Emily Fertig²

Sharad Vikram²

Marcel van Gerven¹

¹Department of Artificial Intelligence, Donders Institute for Brain, Cognition and Behaviour
Radboud University, Nijmegen, the Netherlands

²Google Research, San Francisco, CA, USA

Abstract

Stochastic variational inference offers an attractive option as a default method for differentiable probabilistic programming. However, the performance of the variational approach depends on the choice of an appropriate variational family. Here, we introduce automatic structured variational inference (ASVI), a fully automated method for constructing structured variational families, inspired by the closed-form update in conjugate Bayesian models. These convex-update families incorporate the forward pass of the input probabilistic program and can therefore capture complex statistical dependencies. Convex-update families have the same space and time complexity as the input probabilistic program and are therefore tractable for a very large family of models including both continuous and discrete variables. We validate our automatic variational method on a wide range of low- and high-dimensional inference problems. We find that ASVI provides a clear improvement in performance when compared with other popular approaches such as the mean-field approach and inverse autoregressive flows. We provide an open source implementation of ASVI in TensorFlow Probability.

1 Introduction

Stochastic gradient-based variational methods have recently emerged as a powerful alternative to MCMC for inference in differentiable probabilistic programming languages (Wingate and Weber, 2013; Kucukelbir et al., 2017; Tran et al., 2016; Kucukelbir et al., 2015;

Bingham et al., 2019). This trend is a consequence of the increasing automatization of variational inference (VI) techniques, which evolved from requiring highly mathematically sophisticated and model-specific tools, to generic algorithms that can be applied to a broad class of problems without model-specific derivations (Hoffman et al., 2013; Kingma and Welling, 2013; Hernández-Lobato and Adams, 2015; Ranganath et al., 2014). However, applications of VI usually require the user to specify a parameterized variational family. In general, it is relatively easy to automate the construction of the variational distribution under a mean-field approximation, where the approximate posterior distribution factorizes as a product of univariate distributions. However, while there is a substantial body of model-specific research on structured variational families, few existing methods can be used for automatically constructing an appropriate scalable structured variational approximation for an arbitrarily chosen probabilistic model. Instead, existing methods either ignore most of the prior structure of the model (e.g. ADVI with the multivariate Gaussian distribution (Kucukelbir et al., 2017)) or require strict assumptions such as local conjugacy (e.g. structured stochastic VI (Hoffman and Blei, 2015)). In addition, several of these methods require the use of *ad-hoc* gradient estimators or variational lower bounds (Tran et al., 2015; Ranganath et al., 2016).

In this paper we introduce an automatic procedure for constructing variational approximations that incorporate the structure of the probabilistic model, while being flexible enough to capture the distribution of the observed data. The construction of these variational approximations is fully automatic and the resulting variational distribution has the same time and space complexity as the input probabilistic program. The new family of variational models, which we call the *convex-update variational family*, interpolates the evidence coming from the observed data with the feedforward structure of the model prior. Specifically, the parameters of the posterior distribution of each latent variable are a convex combination of the parameters induced by the parents of that variable, and a term reflecting

the influence of the data. This mimics the evidence update in the expectation parameters of conjugate exponential family models, where this posterior is given in closed form. The convex-update variational family can be trained using standard inference techniques and gradient estimators and can therefore be used as a drop-in replacement of the mean-field approach in automatic differentiation stochastic VI. We refer to this new form of fully automatic inference as *automatic structured variational inference* (ASVI). A full implementation is open-sourced in TensorFlow Probability as [tfp.experimental.build_asvi_surrogate_posterior](#).

2 Preliminaries

VI is used to approximate the posterior over the latent variables of a probabilistic program $p(\mathbf{x})$ with a member of a parameterized family of probability distributions $q(\mathbf{x}; \boldsymbol{\psi})$. The vector of variational parameters $\boldsymbol{\psi}$ is obtained by maximizing the evidence lower bound (ELBO):

$$\mathcal{L}[\boldsymbol{\psi}] = -\mathbb{E}_{q(\mathbf{x}; \boldsymbol{\psi})} \left[\log \frac{q(\mathbf{x}; \boldsymbol{\psi})}{L(\mathbf{y} | \mathbf{x}) p(\mathbf{x})} \right], \quad (1)$$

where $L(\mathbf{y} | \mathbf{x})$ is a likelihood function. The resulting variational posterior (i.e. the maximum of this optimization problem) depends on the choice of the parameterized family $q(\mathbf{x}; \boldsymbol{\psi})$ and is equal to the exact posterior only when the latter is included in the family. In this paper we restrict our attention to probabilistic programs that are specified in terms of conditional probabilities and densities chained together by deterministic functions:

$$p(\mathbf{x}) = \prod_j \rho_j(x_j | \boldsymbol{\theta}_j(\boldsymbol{\pi}_j)), \quad (2)$$

where $\rho_j(\cdot | \cdot)$ is a family of probability distributions and $\boldsymbol{\pi}_j \subseteq \{x_i\}_{i \neq j}$ is a subset of parent variables such that the resulting graphical model is a directed acyclic graph (DAG). The vector-valued functions $\boldsymbol{\theta}_j(\boldsymbol{\pi}_j)$ specify the value of the parameters of the distribution of the latent variable x_j given the values of all its parents.

2.1 Convex updates in conjugate models

Exponential family distributions play a central role in Bayesian statistics as the only families that support conjugate priors, where the posterior is available via an analytic expression (Diaconis and Ylvisaker, 1979). In several important conjugate Bayesian models, the posterior expectation parameters are a convex combination of the prior parameters and the maximum likelihood estimator (Diaconis and Ylvisaker, 1979):

$$\bar{\mu} = \lambda \bar{\mu}_0 + (1 - \lambda) \mu_{\text{ML}}, \quad (3)$$

where λ is a convex combination coefficient. For example, in a Gaussian model with known precision τ and a Gaussian prior over the mean, the posterior mean is

given by

$$\bar{\mu} = \frac{\tau_0}{\tau_0 + N\tau} \bar{\mu}_0 + \frac{N\tau}{\tau_0 + N\tau} \left(\frac{1}{N} \sum_{n=1}^N y_n \right), \quad (4)$$

where τ_0 is the precision of the prior. This formula shows that the posterior parameters are a trade-off between the prior hyper-parameters and the values induced by the data.

3 Convex-update variational families

Consider the following probabilistic model:

$$p(x, y) = L(y | x) \rho(x | \boldsymbol{\theta}), \quad (5)$$

where $L(y | x)$ is a likelihood function and $\rho(x | \boldsymbol{\theta})$ is a prior distribution parameterized by a vector of parameters $\boldsymbol{\theta}$. We do not assume the likelihood or the prior to be in the exponential family; simply that the parameters $\boldsymbol{\theta}$ are defined in a convex domain. We can construct a convex-update parameterized variational family by mimicking the form of the parameter update rule in conjugate models:

$$q(x; \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \rho(x | \boldsymbol{\lambda} \odot \boldsymbol{\theta} + (1 - \boldsymbol{\lambda}) \odot \boldsymbol{\alpha}), \quad (6)$$

where $\boldsymbol{\lambda}$ is now a vector of learnable parameters with entries ranging from 0 to 1 and $\boldsymbol{\alpha}$ is a vector of learnable parameters with the same range of possible values as $\boldsymbol{\theta}$. In practice, it is convenient to express each λ_i as the logistic sigmoid of a logit parameter, so that we can perform unconstrained optimization in the logit space.

In a model with a single latent variable, the convex-update parameterization is over-parameterized and equivalent to mean field. However, the power of this approach becomes evident in multivariate models constructed by chaining basic probability distributions together. Consider a probabilistic program specified in the form of Eq. (2). We can construct a structured variational family by applying the convex-update form to each latent conditional distribution in the model:

$$q(\mathbf{x}; \boldsymbol{\Lambda}, \mathbf{A}) = \prod_j \rho_j(x_j | \mathcal{U}_{\boldsymbol{\lambda}_j}^{\boldsymbol{\alpha}_j}[\boldsymbol{\theta}_j(\boldsymbol{\pi}_j)]) \quad (7)$$

with $\boldsymbol{\Lambda} = (\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_J)$, $\mathbf{A} = (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_J)$ and the *convex update operator*:

$$\mathcal{U}_{\boldsymbol{\lambda}}^{\boldsymbol{\alpha}}[\boldsymbol{\theta}] = \boldsymbol{\lambda} \odot \boldsymbol{\theta} + (1 - \boldsymbol{\lambda}) \odot \boldsymbol{\alpha}. \quad (8)$$

The multivariate structured distributions induced by this family have several appealing theoretical properties that justify their usage in structured inference problems:

1. The family always contains the original probabilistic program (i.e. the prior distribution). This is trivial to see as we can obtain the prior by setting $\boldsymbol{\lambda} = \mathbf{1}$. On the other hand, setting $\boldsymbol{\lambda} = \mathbf{0}$ results in the standard mean-field approximation. Note that none of the commonly used automatic structured variational approaches share this basic property.

```

import tensorflow as tf
import tensorflow_probability as tfp
tfd = tfp.distributions
num_timesteps = 20
innovation_scale = 0.1

@tfd.JointDistributionCoroutineAutoBatched
def brownian_motion():
    xt = 0.
    for t in range(num_timesteps):
        xt = yield tfd.Normal(loc=xt,
                               scale=innovation_scale)
    
```

```

# Omitted for brevity: definitions of variational
# parameters 'lam' and 'alpha' as 'tf.Variable's.
@tfd.JointDistributionCoroutineAutoBatched
def brownian_motion_surrogate_posterior():
    xt = 0.
    for t in range(num_timesteps):
        xt = yield tfd.Normal(
            loc=(lam[t]['loc'] * xt +
                 (1 - lam[t]['loc']) * alpha[t]['loc']),
            scale=(lam[t]['scale'] * innovation_scale +
                  (1 - lam[t]['scale']) *
                  alpha[t]['scale']))
    
```

Figure 1: TensorFlow Probability models for a discretized Brownian motion prior (left) and corresponding structured variational surrogate (right). In TFP’s coroutine syntax, **yielding** a distribution samples a new random variable. The surrogate program is a written-out version of the process implicitly (and automatically) defined by the automatic method in Figure 2. Note that it has the same control-flow structure as the original model.

2. The family includes both the exact filtering and smoothing posterior of univariate linear Gaussian time series models of the form

$$x_t \sim \mathcal{N}(ax_{t-1}, \sigma^2), \quad y_t \sim \mathcal{N}(x_t, \xi^2). \quad (9)$$

In this case, the filtering conditional posterior is given by the Kalman filter update:

$$\begin{aligned} \bar{\mu}_{t+1} &= (1 - K_t)\mu_{t+1}(x_t) + K_ty_t \\ &= (1 - K_t)ax_t + K_ty_t, \end{aligned} \quad (10)$$

where $0 \leq K_t \leq 1$ is the Kalman gain, corresponding to λ_t in ASVI. The smoothing update has a similar form where the data term is augmented with an estimate integrating all future observations.

3. The convex-update family has a very parsimonious parameterization compared to other structured families. The number of parameters is $2P$, where P is the total number of parameters of the conditional distributions. In contrast, the common multivariate normal approach scales as P^2 . However, this parsimonious parameterization implies that the convex-update family cannot capture dependencies that are not already present in the prior of the probabilistic program. Specifically, the convex-update family cannot model correlations originating from colliding arrows in the directed acyclic graph (‘explaining away’ dependencies).

4 Discrete latent variables and stochastic control flow

While this paper is focused on differentiable probabilistic programming, the convex-update family can be applied equally well to models with discrete latent variables and combinatorial link functions. Consider a conditional Bernoulli variable:

$$p(b \mid \pi) = \rho(\pi)^b (1 - \rho(\pi))^{1-b}, \quad (11)$$

where the array π collects the values of its parents. Since the parameter ρ is defined in a convex domain,

we can construct a convex-update variational posterior:

$$p(b \mid \pi) = \mathcal{U}_\lambda^\alpha [\rho(\pi)]^b (1 - \mathcal{U}_\lambda^\alpha [\rho(\pi)])^{1-b}. \quad (12)$$

It is easy to see that the same procedure can be applied to binomial, categorical, multinomial and Poisson variables among others. Using discrete variables, we can also construct convex-update families with stochastic control flows. For example, a binary latent b can determine whether the stochastic computation proceeds on one branch of the computation tree rather than another. This has the nice feature of having an approximate posterior program with the same stochastic flow structure of the prior program where the probability of each gate is biased by the observed data.

5 Automatic structured variational inference

Construction of the variational family (7) can be straightforwardly automated by probabilistic programming systems. A variational program is defined by running the input program under a nonstandard interpretation: at each random variable site

$$x_j^{\text{prior}} \sim p_j(x_j \mid \theta_j(\pi_j^{\text{prior}})),$$

the variational program instead samples from the convex-update distribution (Eq. (6)),

$$x_j^{\text{posterior}} \sim q(x_j \mid \theta_j(\pi_j^{\text{posterior}}); \lambda_j, \alpha_j),$$

where λ_j and α_j are trainable variables. Because each variable definition is rewritten ‘locally’, the variational program has the same control-flow structure, dependence graph, and time complexity as the original model. Local transformations of this kind may be implemented by effect handlers (Plotkin and Pretnar, 2009), a mechanism supported in multiple recent probabilistic programming frameworks (Moore and Gorinova, 2018; Bingham et al., 2019; Phan et al., 2019).

We provide an open-source implementation of ASVI for joint distribution models in TensorFlow Probabil-

```

def build_asvi_surrogate_posterior(prior):
    ConvexUpdateParams = namedtuple(['lam', 'alpha'])
    q_vars = {}
    @tf.nn.joint_distribution_coroutine_auto_batched
    def asvi_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] = {
                    k: ConvexUpdateParams(
                        lam=Variable(random_init_like(v),
                                    xform=Sigmoid()),
                        alpha=Variable(random_init_like(v),
                                      xform=dist.bijectors[k]()))
                    for k, v in rv.parameters.items()}
            # Apply convex update for each parameter.
            q_params = {}
            for k, v in rv.parameters.items():
                lam = q_vars[rv.name][k].lam
                alpha = q_vars[rv.name][k].alpha
                q_params[k] = lam * v + (1. - lam) * alpha
            # Step to the next RV.
            q_sample = yield type(rv)(**q_params)
            rv = prior_gen.send(q_sample)
    return asvi_surrogate_posterior

```

Figure 2: Python pseudocode for general-purpose automated construction of a structured variational family. A full implementation is provided in TensorFlow Probability ([tfp.experimental.build_asvi_surrogate_posterior](#)).

ity (TFP; Dillon et al., 2017), which uses an effect-handling-like mechanism (Python coroutines) to transform an input joint distribution into a structured variational surrogate. Figure 2 shows a simplified version.

Our implementation operates in the default parameterization θ of each TFP distribution class, which is typically similar (if not exactly equal) to the expectation parameterization; e.g., normal distributions expose a location-scale parameterization. Constrained variational parameters (λ_j, α_j) are defined as differentiable transformations of unconstrained variables, where each distribution is annotated with appropriate constraining transformations for its parameters. For example, the softplus transformation $f(x) = \log(1 + \exp(x))$ enforces that scale parameters have positive values. For variables such as log-normal distributions whose definition involves a bijective transformation $\mathbf{y} = f_\psi(\mathbf{x})$ of a base variable $\mathbf{x} \sim p_\phi(\mathbf{x})$, the convex-update family has an analogous form and defines a convex update on the union of the parameters $\theta = (\phi, \psi)$.

Given a structured variational program, the ELBO (1) may be estimated by Monte Carlo sampling and optimized with stochastic gradient estimates, obtained via the reparameterization trick for continuous-valued variables and score-function estimator for discrete variables (Ranganath et al., 2014; Kucukelbir et al., 2017).

6 Related work

Most structured VI approaches require model specific derivations and variational bounds. However, several forms of model-agnostic structured variational distributions have been introduced. The most commonly used fully automatic approach is probably automatic differentiation variational inference (ADVI) (Kucukelbir et al., 2017). The ADVI variational family is constructed by mapping the values of all latent variables to an unbounded coordinate space based on the support of each distribution. The variational distribution in this new space is then parameterized as either a spherical (mean field) or a fully multivariate normal distribution. While this approach is broadly applicable, it exploits very little information from the original probabilistic model and has scalability problems due to the cubic complexity of Bayesian inference when the multivariate distribution is used. Hierarchical VI accounts for dependencies between latent variables by coupling the parameters of their factorized distributions through a joint *variational prior* (Ranganath et al., 2016). While this method is very generic, it requires user input in order to define the variational prior and it uses a modified variational lower bound. Copula VI models the dependencies between latent variables using a vine copula function (Tran et al., 2015). In the context of probabilistic programming, copula VI shares some of the same limitations of hierarchical VI: it requires the appropriate specification of bivariate copulas and it needs a specialized inference technique. The approach that is closest to our current work is structured stochastic VI (Hoffman et al., 2013). Similar to our approach, its variational posteriors have the same conditional independence structure as the input probabilistic program. However, this method is limited to conditionally conjugate models with exponential family distributions. On the other hand, conjugate-Computation variational inference can be applied to general non-conjugate models and the structure of the optimal distribution is identical to the structure of the model. (Khan and Lin, 2017).

Normalizing flows are general deep learning methods designed to approximate arbitrary probability densities through a series of learnable invertible mappings with tractable Jacobians (Rezende and Mohamed, 2015; Kingma et al., 2016; Dinh et al., 2016; Kingma and Dhariwal, 2018; Papamakarios et al., 2017; Kobyzev et al., 2020). Flows have been used in stochastic VI as very expressive and highly parameterized variational families (Rezende and Mohamed, 2015; Kingma et al., 2016). Most normalizing flow architectures used in VI do not incorporate the conditional independence structure of the prior probabilistic program. Structured conditional continuous normalizing flows are a new class of normalizing flows that have the conditional independence structure of the true posterior (Weilbach et al., 2020). These architectures are based on faithful

inversion (Webb et al., 2018) and implement the dependency structure using sparse matrices. Conversely to ASVI, this approach only exploits the graphical structure of the probabilistic program and therefore ignores the specific form of the conditional distributions and link functions.

Structured VI is most commonly applied in time series models such as hidden Markov models and autoregressive models. In these cases the posterior distributions inherit strong statistical dependencies from the sequential nature of the prior. Structured VI for time series usually use structured variational families that capture the temporal dependencies, while being fully-factorized in the non-temporal variables (Eddy, 1996; Foti et al., 2014; Johnson and Willsky, 2014; Karl et al., 2016; Fortunato et al., 2017). This differs from our convex-update families, where both temporal and non-temporal dependencies are preserved.

7 Applications

We first evaluate ASVI on a set of standardized Bayesian inference tasks to compare ASVI to other automated inference methods. We then apply ASVI to a both a deep Bayesian smoothing and a deep generative modeling task, to demonstrate that ASVI scales to large datasets while producing high quality posteriors. The code needed to run ASVI and baselines on the standardized tasks can be found in [this repository](#).

7.1 Inference Gym tasks

Inference Gym (IG; Sountsov et al., 2020) is a Python library for evaluating probabilistic inference algorithms. It defines a set of standardized inference tasks and provides implementations as TFP joint distribution models (Piponi et al., 2020) and as Stan programs (Carpenter et al., 2017), which are used to compute ‘ground truth’ posterior statistics via MCMC.

For each task we evaluate ASVI against a suite of baseline variational posteriors, including mean field (MF), inverse autoregressive flows (Kingma et al., 2016; Papamakarios et al., 2017) with eight hidden units (Small IAF) and 512 hidden units (Large IAF), and a multivariate normal posterior (MVN). Where possible, we also compare to an AR(1) posterior. For each method the ELBO is optimized to convergence (up to 100000 steps, although significantly fewer were required in most cases) using the Adam optimizer, with learning rate set by a hyperparameter sweep for each task/method pair—this tuning was particularly important for the IAF models. Further details are provided in the supplement.

Time series. We consider two discretized stochastic differential equation (SDE) models, defined by condi-

tional densities of the following form:

$$x_{t+1} \sim \mathcal{N}(x_t + f(x_t, t) dt, g(x_t, t)^2 dt) \quad (13)$$

$$y_t \sim \mathcal{N}(x_t, \sigma_{\text{obs}}^2) \quad (14)$$

where x represents the latent process with drift function f and volatility function g , and observations y are taken with noise standard deviation σ_{obs} . The first model describes Brownian (BR) motion without drift, and the second model is a stochastic Lorenz dynamical system (LZ). Full model specifications are provided in the supplementary material. We also include variants of both models that include global variables (BRG and LZG) where the innovation and observation noise scale parameters are unknown. We simulate each model forward 30 steps and assume that noisy observations are taken of the first ten and last ten steps, and the middle ten are unobserved.

Hierarchical regression models. We also evaluate ASVI on two standard Bayesian hierarchical models: Eight Schools (ES; Gelman et al., 2013), which models the effect of coaching programs on standardized test scores, and Radon (R; Gelman and Hill, 2007), a hierarchical linear regression model that predicts Radon measurements taken in houses. Full descriptions are provided in the supplementary material.

Results. Table 1 shows that ASVI and the large IAF produce the best ELBOs; ASVI is competitive with the large IAF on all tasks despite having far fewer parameters. Similarly, in Table 2 we see that ASVI’s estimates of posterior means and standard deviations are consistently among the best relative to MCMC ground truth. Mean field fails to converge to the true posterior estimate in all four models, indicating that inclusion of prior structure in ASVI is needed to capture the dependencies in the true posterior. Qualitatively, we observe in Figure 3 that only the ASVI posterior captures the ground truth posterior in the Lorenz task, indicating that ASVI is able to capitalize on strong prior structure when it exists. Figure 1 in the supplementary material shows loss trajectories. For all but the Radon model, ASVI tends to converge in fewer iterations than the large IAF, the other highest-performing surrogate posterior.

7.2 Deep Bayesian smoothing with neural SDEs

So far, we performed inference in simple time series and hierarchical models with relatively low-dimensional state spaces. We will now test the performance of ASVI on a high-dimensional problem with complex nonlinearities parameterized by deep networks. As a latent model, we use neural (discretized) stochastic differential equations (SDE) (Chen et al., 2018; Li et al., 2020):

$$d\mathbf{x} = \mathbf{F}(\mathbf{x})dt + d\mathbf{B}(t) \quad (15)$$

where $\mathbf{F}(\mathbf{x})$ is a nonlinear function parameterized by a neural network and $\mathbf{B}(t)$ is standard multivariate

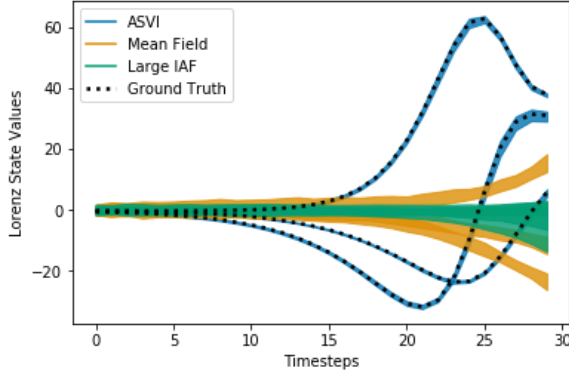


Figure 3: Qualitative results of the Lorenz system with globals (LZG) experiment, showing each baseline’s estimate of the three latent dimensions in the Lorenz model \pm two standard deviations. Only ASVI is able to successfully infer the values of each latent series.

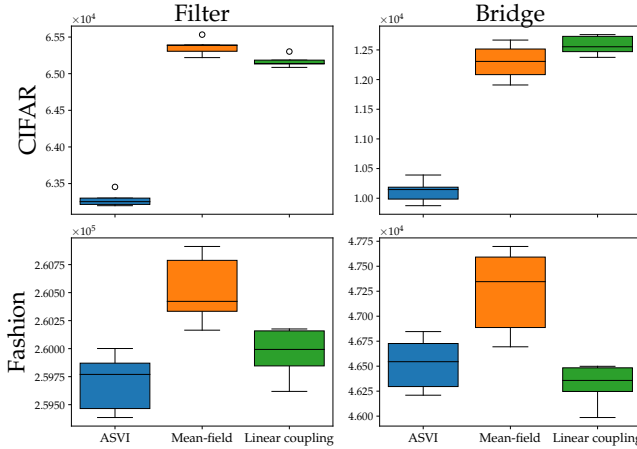


Figure 4: Quantitative results of the deep Bayesian smoothing experiment in subsection 7.2. Shown is the negative ELBO of ASVI, mean-field, and linear coupling at the final training iteration.

Brownian motion (see Supplementary Material C for the details of the architecture). This latent process generates noise-corrupted observations through a deep network \mathbf{G} :

$$\mathbf{y}_t \sim \mathcal{N}(\mathbf{G}(\mathbf{x}(t)), \sigma^2). \quad (16)$$

In our examples, \mathbf{G} is a generator which converts latent vectors into RGB images. We tested two kinds of pre-trained generators: A DCGAN trained on CIFAR10 and a DCGAN trained on FashionGEN (Radford et al., 2015) (see Supplementary Material C). In the former case, the latent space is 100-dimensional, while in the latter it is 120-dimensional. We considered two inference challenges: filtering, where we aim to remove the noise from a series of images generated by a trajectory in the latent process, and bridging, where we reconstruct a series of intermediate observations given the beginning (first three time points)

and the end (last two time points) of a trajectory. In both cases, we assume knowledge of the dynamical and generative models, discretize the neural SDE using an Euler–Maruyama scheme and backpropagate through the integrator (Chen et al., 2018).

Figure 5 shows the filtering performance of ASVI and two baselines (mean field, and linear Gaussian model; see Supplementary Material C for the details) in a bridging problem. The quantitative results (negative ELBOs) are shown in Figure 4. ASVI always reaches tighter lower bounds except in the Fashion bridge experiment where it has slightly lower performance than the linear coupling baseline. This tighter variational bound results in discernibly higher quality filtered images in both CIFAR10 and Fashion, as shown in Figure 5. The figure also shows several samples from the ASVI bridge posterior. As expected, the generated images diverge in the unobserved period and reconverge at the end.

7.3 Deep amortized generative modeling

Finally, we apply an amortized form of ASVI to a deep generative modeling problem. The goal is to model the joint distribution of a set of binary images \mathbf{y} paired with class labels l . We use a deep variational autoencoder with three layers of latent variables \mathbf{z}_1 , \mathbf{z}_2 and \mathbf{z}_3 coupled in a feed-forward fashion through ReLU fully-connected networks:

$$\begin{aligned} \mathbf{z}_1 &\sim \mathcal{N}(0, 1.5) \\ \mathbf{z}_2 &\sim \mathcal{N}(\text{ReLU}(\mathbf{f}_1(\mathbf{z}_1)), 0.25) \\ \mathbf{z}_3 &\sim \mathcal{N}(\text{ReLU}(\mathbf{f}_2(\mathbf{z}_2)), 0.1) \\ \mathbf{y} &\sim \mathcal{N}(\mathbf{g}_1(\mathbf{z}_3), 0.1) \\ l &\sim \text{Categorical}(\text{Softmax}(\mathbf{g}_2(\mathbf{z}_2))) \end{aligned} \quad (17)$$

where \mathbf{f}_1 and \mathbf{f}_2 are fully-connected two-layer networks with ReLU activations, linear output units, and 25 and 75 hidden units respectively, while \mathbf{g}_1 and \mathbf{g}_2 are linear layers. The amortized convex-update distribution has the following form:

$$\begin{aligned} \mathbf{z}_1 &\sim \mathcal{N}(\boldsymbol{\alpha}_1(\mathbf{y}), \boldsymbol{\xi}_1(\mathbf{y})) \\ \mathbf{z}_2 &\sim \mathcal{N}(\boldsymbol{\lambda}_2 \mathbf{f}_1(\mathbf{z}_1) + (1 - \boldsymbol{\lambda}_2) \boldsymbol{\alpha}_2(\mathbf{y}), \boldsymbol{\xi}_2(\mathbf{y})) \\ \mathbf{z}_3 &\sim \mathcal{N}(\boldsymbol{\lambda}_3 \mathbf{f}_2(\mathbf{z}_1) + (1 - \boldsymbol{\lambda}_3) \boldsymbol{\alpha}_3(\mathbf{y}), \boldsymbol{\xi}_3(\mathbf{y})) \end{aligned} \quad (18)$$

where the mean vectors $\boldsymbol{\alpha}_k(\mathbf{y})$ are the activation of the $(8 - 2k)$ -th layer (post ReLU) of a fully-connected 6-layer ReLU inference network (with layer sizes (120, 100, 70, 50, 70, 2)) that takes the image \mathbf{y} as input. The scale parameter vectors $\boldsymbol{\xi}_k(\mathbf{y})$ were obtained by connecting a linear layer to the $(8 - 2k)$ -th layer, followed by a softplus transformation. The details of all architectures are given in Supplementary Material D. The amortized family was parameterized by $\boldsymbol{\lambda}$ and the weights and biases of the inference network. Note that this amortization scheme is not fully automatic as it requires the choice of the inference networks. The mean-field baseline had the same form given in Eq. (18), but with the expectation of the distribution fully determined by $\boldsymbol{\alpha}_k(\mathbf{y})$. We did not include a comparison with

Table 1: Average final negative ELBO values and their standard errors, computed over 15 simulations. Each row corresponds to an Inference Gym task and each column is a posterior baseline. Boldface indicates the best performance.

	ASVI	Mean field	Small IAF	Large IAF	MVN	AR(1)
BR	-5.20 ± 0.11	0.99 ± 0.49	-3.47 ± 0.24	-4.9 ± 0.12	-3.7 ± 0.2	-4.01 ± 0.31
BRG	-0.35 ± 0.10	-5.16 ± 0.31	-1.31 ± 0.23	0.14 ± 0.09	-2.27 ± 0.20	-2.08 ± 0.31
LZ	34.29 ± 0.66	1225.66 ± 4.30	1236.49 ± 12.47	1242.11 ± 10.39	2153.08 ± 179.24	1271.36 ± 16.53
LZG	34.68 ± 0.15	119.67 ± 0.17	112.64 ± 0.58	94.8 ± 0.14	98.43 ± 0.09	54.53 ± 2.1
ES	36.50 ± 0.04	36.94 ± 0.04	36.28 ± 0.04	36.26 ± 0.02	36.58 ± 0.02	N/A
R	1079.8 ± 0.22	1082.77 ± 0.59	1078.85 ± 0.21	1078.52 ± 0.19	1090.96 ± 0.48	N/A

Table 2: Mean and standard error of the absolute error in posterior mean (M) and standard deviation (SD) parameters relative to MCMC ground truth, normalized by true SD, computed over 15 simulations. Boldface indicates the best performance.

		ASVI	Mean field	Small IAF	Large IAF	MVN	AR(1)
BR	M	0.16 ± 0.03	0.14 ± 0.02	0.1 ± 0.02	0.12 ± 0.03	0.09 ± 0.01	0.24 ± 0.05
	SD	0.06 ± 0.02	0.37 ± 0.01	0.14 ± 0.03	0.07 ± 0.02	0.07 ± 0.01	0.07 ± 0.01
BRG	M	0.69 ± 0.07	0.81 ± 0.03	0.77 ± 0.03	0.71 ± 0.05	0.74 ± 0.03	0.80 ± 0.07
	SD	0.22 ± 0.02	0.31 ± 0.01	0.27 ± 0.02	0.23 ± 0.02	0.23 ± 0.02	0.24 ± 0.02
LZ	M	0.36 ± 0.25	35.83 ± 0.00	35.94 ± 0.11	36.62 ± 0.07	39.66 ± 0.00	36.52 ± 0.18
	SD	0.47 ± 0.03	0.94 ± 0.00	0.92 ± 0.01	0.81 ± 0.01	0.67 ± 0.01	0.83 ± 0.01
LZG	M	0.15 ± 0.03	21.46 ± 0.04	23.35 ± 0.16	23.96 ± 0.06	23.86 ± 0.08	1.56 ± 0.24
	SD	0.39 ± 0.02	1.51 ± 0.04	1.51 ± 0.20	0.80 ± 0.14	2.59 ± 0.06	2.59 ± 0.08
ES	M	0.16 ± 0.03	0.16 ± 0.02	0.16 ± 0.04	0.13 ± 0.04	0.13 ± 0.04	N/A
	SD	0.07 ± 0.02	0.05 ± 0.01	0.07 ± 0.02	0.08 ± 0.03	0.06 ± 0.01	N/A
R	M	0.12 ± 0.02	0.24 ± 0.02	0.17 ± 0.02	0.15 ± 0.02	0.42 ± 0.01	N/A
	SD	0.15 ± 0.03	0.12 ± 0.01	0.09 ± 0.01	0.10 ± 0.02	0.50 ± 0.02	N/A




Figure 5: Qualitative results of the deep Bayesian filters using ASVI, mean-field, and linear coupling for CIFAR10 (left) and FashionGEN (right). The top row shows the ground truth images, the second row the noisy-corrupted observations, and the remaining rows show the different reconstructions.

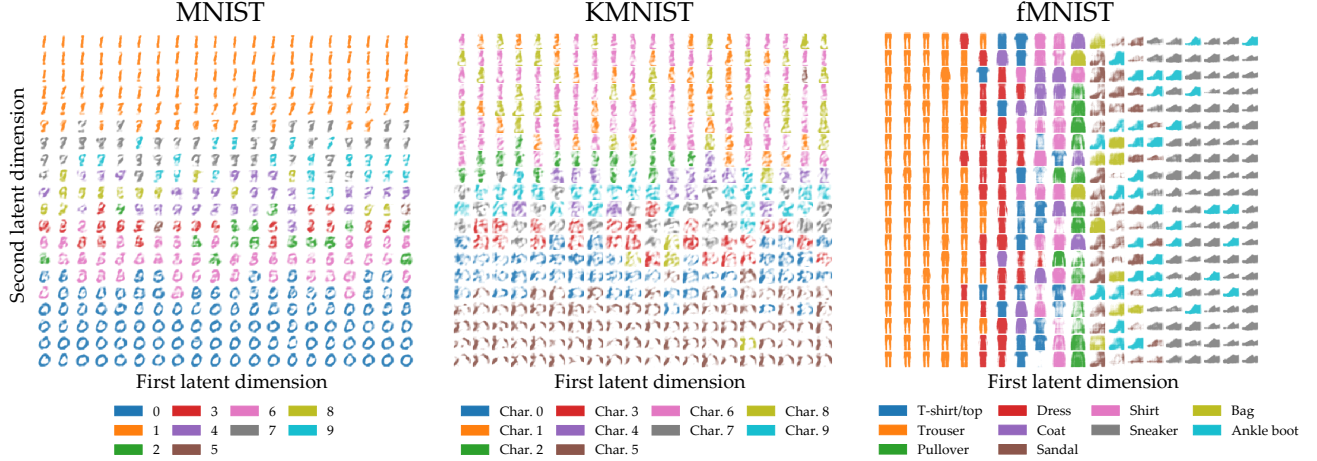


Figure 6: The latent embeddings for the different MNIST variants. Each element is a predicted sample together with its corresponding label.

Table 3: Average and standard error of the RMSE between the posterior means and the ground truth functions, computed over 15 simulations. Boldface indicates the best performance.

	MNIST		kMNIST		fMNIST	
	no labels	labels	no labels	labels	no labels	labels
ASVI	-134.6 \pm 0.1	-99.7 \pm 0.5	-252.4 \pm 0.3	-200.3 \pm 0.7	-151.1 \pm 0.4	-116.4 \pm 1.2
MF	-153.8 \pm 0.2	-146.3 \pm 0.5	-267.0 \pm 0.3	-258.6 \pm 0.8	-166.7 \pm 0.2	-158.2 \pm 1.3

the other baselines as it is computationally unfeasible in this larger scale experiment.

We tested the performance of these deep variational generative models on three computer vision datasets: MNIST, FashionMNIST and KMNIST (LeCun et al., 1998; Xiao et al., 2017; Clanuwat et al., 2018). Furthermore, we performed two types of experiment. In the first, images and labels were generated jointly. In the second, only images were generated. Table 3 reports the performance of ASVI and mean-field baseline quantified as the test set ELBO. The ELBO was estimated by sampling the latent variables 20 times per test image from the variational posterior. Mean ELBO and SEM were obtained by repeating the estimation 20 times. ASVI achieves higher performance in all experiments for all datasets. Figure 6 shows a randomized selection of images generated by the ASVI model together with the corresponding label.

8 Discussion

We introduced an automatic algorithm for constructing an appropriately structured variational family, given an input probabilistic program. The resulting method can be used on any probabilistic program specified by a directed Bayesian network, and always preserves the forward-pass structure of the input program. The main limitation of the convex-update family is that it cannot capture dependencies induced by colliding arrows in

the input graphical model. Consequently, in a model such as a standard Bayesian neural network, where the prior over the weights is decoupled, the convex-update family is a mean-field family. Despite this limitation, our results demonstrate good performance (competitive with a high-capacity IAF) on small Bayesian hierarchical models with colliding arrows in the prior graphical model, through capturing the forward structure alone. Designing structured surrogates that efficiently capture the full posterior dependence structure is an interesting direction for ongoing work.

References

- D. Wingate and T. Weber. Automated variational inference in probabilistic programming. *ArXiv preprint arXiv:1301.1299*, 2013.
- 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.
- 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.
- A. Kucukelbir, R. Ranganath, A. Gelman, and D. Blei. Automatic variational inference in Stan. In *Advances in Neural Information Processing Systems*, 2015.
- 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.
- M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- D. P. Kingma and M. Welling. Auto-encoding variational Bayes. *ArXiv preprint arXiv:1312.6114*, 2013.
- J. M. Hernández-Lobato and R. Adams. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *International Conference on Machine Learning*, pages 1861–1869, 2015.
- R. Ranganath, S. Gerrish, and D. Blei. Black box variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2014.
- M. Hoffman and D. Blei. Stochastic structured variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2015.
- D. Tran, D. Blei, and E. M. Airoldi. Copula variational inference. In *Advances in Neural Information Processing Systems*, 2015.
- R. Ranganath, D. Tran, and D. Blei. Hierarchical variational models. In *International Conference on Machine Learning*, 2016.
- P. Diaconis and D. Ylvisaker. Conjugate priors for exponential families. *The Annals of Statistics*, 7(2): 269–281, 1979.
- G. Plotkin and M. Pretnar. Handlers of algebraic effects. In *European Symposium on Programming*. Springer, 2009.
- D. Moore and M. I. Gorinova. Effect handling for composable program transformations in Edward2. *arXiv preprint arXiv:1811.06150*, 2018.
- D. Phan, N. Pradhan, and M. Jankowiak. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *ArXiv preprint arXiv:1912.11554*, 2019.
- 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.
- M. E. Khan and W. Lin. Conjugate-computation variational inference: Converting variational inference in non-conjugate models to inferences in conjugate models. In *International Conference on Artificial Intelligence and Statistics*, 2017.
- D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.
- 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*, 2016.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real NVP. *arXiv preprint arXiv:1605.08803*, 2016.
- D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, 2018.
- G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2017.
- I. Kobyzev, S. Prince, and M. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- 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*, 2020.
- S. Webb, A. Golinski, R. Zinkov, N. Siddharth, T. Rainforth, Y. W. Teh, and F. Wood. Faithful inversion of generative models for effective amortized inference. In *Advances in Neural Information Processing Systems*, 2018.
- 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 *Advances 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.
- P. Sountsov, A. Radul, and contributors. Inference gym, 2020.
- D. Pisoni, D. Moore, and J. V. Dillon. Joint distributions for TensorFlow Probability. *ArXiv preprint arXiv:2001.11819*, 2020.
- B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.
- A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*. CRC Press, 3rd edition, 2013.
- A. Gelman and J. Hill. *Data Analysis using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, 2007.
- T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, 2018.
- X. Li, T. L. Wong, R. T.Q. Chen, and D. Duvenaud. Scalable gradients for stochastic differential equations. *arXiv preprint arXiv:2001.01328*, 2020.
- A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *ArXiv preprint arXiv:1511.06434*, 2015.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep learning for classical Japanese literature. *ArXiv preprint arXiv:1812.01718*, 2018.