



HAL
open science

Monte Carlo Variational Auto-Encoders

Achille Thin, Nikita Kotelevskii, Alain Durmus, Maxim Panov, Eric Moulines,
Arnaud Doucet

► **To cite this version:**

Achille Thin, Nikita Kotelevskii, Alain Durmus, Maxim Panov, Eric Moulines, et al.. Monte Carlo Variational Auto-Encoders. 38th International Conference on Machine Learning, Aug 2021, Online, France. 10.48550/arXiv.2106.15921 . hal-04399045

HAL Id: hal-04399045

<https://hal.science/hal-04399045v1>

Submitted on 12 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Monte Carlo Variational Auto-Encoders

Achille Thin¹ Nikita Kotelevskii² Alain Durmus³ Maxim Panov² Eric Moulines^{1,4,5} Arnaud Doucet⁶

Abstract

Variational auto-encoders (VAE) are popular deep latent variable models which are trained by maximizing an Evidence Lower Bound (ELBO). To obtain tighter ELBO and hence better variational approximations, it has been proposed to use importance sampling to get a lower variance estimate of the evidence. However, importance sampling is known to perform poorly in high dimensions. While it has been suggested many times in the literature to use more sophisticated algorithms such as Annealed Importance Sampling (AIS) and its Sequential Importance Sampling (SIS) extensions, the potential benefits brought by these advanced techniques have never been realized for VAE: the AIS estimate cannot be easily differentiated, while SIS requires the specification of carefully chosen backward Markov kernels. In this paper, we address both issues and demonstrate the performance of the resulting Monte Carlo VAEs on a variety of applications.

1. Introduction

Variational Auto-Encoders (VAE) introduced by (Kingma & Welling, 2013) are a very popular class of methods in unsupervised learning and generative modelling. These methods aim at finding a parameter θ maximizing the marginal log-likelihood $p_\theta(x) = \int p_\theta(x, z) dz$ where $x \in \mathbb{R}^N$ is the observation and $z \in \mathbb{R}^d$ is the latent variable. They rely on the introduction of an additional parameter ϕ and a family of variational distributions $q_\phi(z|x)$. The joint parameters $\{\theta, \phi\}$ are then inferred through the optimization of the

^{*}Equal contribution ¹CMAP, Ecole Polytechnique, Université Paris-Saclay, France ²CDISE, Skolkovo Institute of Science and Technology, Moscow, Russia ³Ecole Nationale Supérieure Paris-Saclay, France ⁴HDI Lab, HSE University, Moscow, Russia ⁵Centre de recherche Lagrange en mathématiques et calcul ⁶University of Oxford. Correspondence to: Achille Thin <achille.thin@polytechnique.edu>.

Evidence Lower Bound (ELBO) defined as

$$\begin{aligned} \mathcal{L}(\theta, \phi) &= \int \log \left(\frac{p_\theta(x, z)}{q_\phi(z|x)} \right) q_\phi(z|x) dz \\ &= \log p_\theta(x) - \text{KL}(q_\phi(z|x) \parallel p_\theta(z|x)) \leq \log p_\theta(x). \end{aligned}$$

The design of expressive variational families has been the topic of many works and is a core ingredient in the efficiency of VAE (Rezende & Mohamed, 2015; Kingma et al., 2016). Another line of research consists in using positive unbiased estimators $\hat{p}_\theta(x)$ of the loglikelihood $p_\theta(x)$ for q_ϕ , *i.e.* $\mathbb{E}_{q_\phi}[\hat{p}_\theta(x)] = p_\theta(x)$. Indeed, as noted in (Mnih & Rezende, 2016), it follows from Jensen's inequality that

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi}[\log \hat{p}_\theta(x)] \leq \log p_\theta(x). \quad (1)$$

A Taylor expansion shows that

$$\mathcal{L}(\theta, \phi) \approx \log p_\theta(x) - \frac{1}{2} \text{var}_{q_\phi} \left[\frac{\hat{p}_\theta(x)}{p_\theta(x)} \right];$$

see e.g. (Maddison et al., 2017; Domke & Sheldon, 2018) for formal results. Hence the ELBO becomes tighter as the variance of the estimator decreases.

A common method to obtain an unbiased estimate is built on importance sampling; *i.e.* $\hat{p}_\theta(x) = n^{-1} \sum_{i=1}^n [p_\theta(x, z_i)/q_\phi(z_i|x)]$ for $z_i \stackrel{\text{i.i.d.}}{\sim} q_\phi(\cdot|x)$. In particular, combined with (1), we obtain the popular Importance Weighted Auto Encoder (IWAE) proposed by (Burda et al., 2015). However, it is expected that the relative variance of this importance-sampling based estimator typically increases with the dimension of the latent z . To circumvent this issue, we suggest in this paper to consider other estimators of the evidence which have shown great success in the Monte Carlo literature. In particular, Annealed Importance Sampling (AIS) (Neal, 2001; Wu et al., 2016), and its Sequential Importance Sampling (SIS) extensions (Del Moral et al., 2006) define state-of-the-art estimators of the evidence. These algorithms rely on an extended target distribution for which an efficient importance distribution can be defined using non-homogeneous Markov kernels.

It has been suggested in various contributions that AIS could be useful to train VAE (Salimans et al., 2015; Wu et al., 2016; Maddison et al., 2017; Wu et al., 2020). However, to the authors knowledge, no contribution discusses how an

unbiased gradient of the resulting ELBO could be obtained. Indeed, the main difficulty in this computation arises from the MCMC transitions in AIS. As a result, when this estimator is used, alternatives to the ELBO (1) have often been considered, see e.g. (Ding & Freedman, 2019).

Whereas AIS requires using MCMC transition kernels, SIS is more flexible and can exploit Markov transition kernels which are only approximately invariant w.r.t. to a given target distribution, e.g. unadjusted Langevin kernels. In this case, the construction of the augmented target distribution, which is at the core of the estimator, requires the careful selection of a class of auxiliary ‘backward’ Markov kernels. (Salimans et al., 2015; Ranganath et al., 2016; Maaløe et al., 2016; Goyal et al., 2017; Huang et al., 2018) propose to learn such auxiliary kernels parameterized with neural networks through the ELBO. However, as illustrated in our simulations, this comes at an increase in the overall computational cost.

Our contributions. The contributions of this paper are as follows:

- (i) We show how to obtain an SIS-based ELBO relying on unadjusted Langevin dynamics which, contrary to (Salimans et al., 2015; Goyal et al., 2017; Huang et al., 2018), does not require introducing and optimizing backward Markov kernels. In addition, an unbiased gradient estimate of the resulting ELBO which exploits the reparameterization trick is derived.
- (ii) We propose an unbiased gradient estimate of the ELBO computed based on AIS. At the core of this estimate is a non-standard representation of Metropolis–Hastings type kernels which allows us to differentiate them. This is combined to a variance reduction technique for improved efficiency.
- (iii) We apply these new methods to build novel Monte Carlo VAEs, and show their efficiency on real-world datasets.

All the theoretical results are detailed in the supplementary material.

2. Variational Inference via Sequential Importance Sampling

2.1. SIS estimator

The design of efficient proposal importance distributions has been proposed in (Crooks, 1998; Neal, 2001) starting from an annealing schedule, and was later extended by (Del Moral et al., 2006). Let $\{\gamma_k\}_{k=0}^K$ be a sequence of unnormalized ‘‘bridge’’ densities satisfying $\gamma_0(z) = q_\phi(z|x)$, $\gamma_K(z) =$

$p_\theta(x, z)$. Note that γ_K is not normalized, in contrast to γ_0 . Here we set

$$\gamma_k(z) = q_\phi(z|x)^{1-\beta_k} p_\theta(x, z)^{\beta_k} \quad (2)$$

for an annealing schedule $0 = \beta_0 < \dots < \beta_K = 1$, but alternative sequences of intermediate densities could be considered. At each iteration, SIS samples a non-homogeneous Markov chain using the transition kernels $\{M_k\}_{k=1}^K$. In this section, we assume that M_k admits a positive transition density m_k , such that m_k leaves γ_k invariant, *i.e.* $\int \gamma_k(z) m_k(z, z') dz = \gamma_k(z')$, or is only approximately invariant. In particular, m_k typically depends on the data x . However, to simplify notation, this dependence is omitted.

Based on this sequence of transition densities, SIS considers the joint density on the path space $z_{0:K} \in \mathbb{R}^{d(K+1)}$

$$q_\phi^K(z_{0:K}|x) = q_\phi(z_0|x) \prod_{k=1}^K m_k(z_{k-1}, z_k), \quad (3)$$

where $z_{i:j} = (z_i, \dots, z_j)$ for $0 \leq i \leq j$. By construction, we expect the K -th marginal $q_\phi^K(z_K|x) = \int q_\phi^K(z_{0:K}|x) dz_{0:K-1}$ to be close to $p_\theta(z_K|x)$. However, we cannot use importance sampling to correct for the discrepancy between these two densities, as $q_\phi^K(z_K|x)$ is typically intractable. To bypass this problem, we introduce another density on the path space

$$p_\theta^K(x, z_{0:K}) = p_\theta(x, z_K) \prod_{k=K-1}^0 \ell_k(z_{k+1}, z_k), \quad (4)$$

where $\{\ell_k\}_{k=0}^{K-1}$ is a sequence of auxiliary positive transition densities. Note that in this case, the K -th marginal $\int p_\theta^K(x, z_{0:K}) dz_{0:K-1}$ is exactly $p_\theta(x, z_K)$. Using now importance sampling on the path space, we obtain the following unbiased SIS estimator (Del Moral et al., 2006; Salimans et al., 2015) by sampling independently $z_{0:K}^i \sim q_\phi^K(\cdot|x)$ and computing

$$\hat{p}_\theta(x) = \frac{1}{n} \sum_{i=1}^n w^K(z_{0:K}^i), \quad w^K(z_{0:K}) = \frac{p_\theta^K(x, z_{0:K})}{q_\phi^K(z_{0:K}|x)}. \quad (5)$$

2.2. AIS estimator

The selection of the kernel $\{\ell_k\}_{k=0}^{K-1}$ has a large impact on the variance of the estimator. The optimal reverse kernels ℓ_k minimizing the variance of $\hat{p}_\theta(x)$ are given by

$$\ell_{k-1}(z_k, z_{k-1}) = \frac{q_{\phi, k-1}(z_{k-1}|x) m_k(z_{k-1}, z_k)}{q_{\phi, k}(z_k|x)}, \quad (6)$$

where $q_{\phi, k}(z_k|x) = \int q_\phi^K(z_{0:K}|x) dz_{0:K}^{-k}$ is the k -th marginal of q_ϕ^K ; see (Del Moral et al., 2006). However,

the resulting estimator $\hat{p}_\theta(x)$ is usually intractable. An approximation to (6) leading to a tractable estimator is provided by AIS (Crooks, 1998; Neal, 2001). When m_k is γ_k -invariant, then, by assuming that $q_{\phi,k} \approx q_{\phi,k-1} \approx \gamma_k$ in (6), we obtain

$$\ell_{k-1}(z_k, z_{k-1}) = \frac{\gamma_k(z_{k-1})m_k(z_{k-1}, z_k)}{\gamma_k(z_k)}. \quad (7)$$

We refer to this kernel as the *reversal* of m_k . In particular, if m_k is γ_k -reversible, *i.e.* $\gamma_k(z_{k-1})m_k(z_{k-1}, z_k) = \gamma_k(z_k)m_k(z_k, z_{k-1})$, then $\ell_{k-1} = m_k$. Note that the weights (5) can be computed using the decomposition $w^K(z_{0:K}) = \prod_{k=1}^K w_k(z_{k-1}, z_k)$ with

$$w_k(z_{k-1}, z_k) = \frac{\gamma_k(z_k)\ell_{k-1}(z_k, z_{k-1})}{\gamma_{k-1}(z_{k-1})m_k(z_{k-1}, z_k)}, \quad (8)$$

which simplifies when using (7) as

$$w_k(z_{k-1}, z_k) = \gamma_k(z_{k-1})/\gamma_{k-1}(z_{k-1}). \quad (9)$$

In contrast to previous works, we will consider in the next section transition kernels m_k which are only approximately γ_k -invariant but still build ℓ_{k-1} in the spirit of (7).

2.3. SIS-ELBO using unadjusted Langevin

For any k , we consider the Langevin dynamics associated to γ_k and the corresponding Euler-Maruyama discretization. Then, we choose for m_k the transition density associated to this discretization

$$m_k(z_{k-1}, z_k) = \mathcal{N}(z_k; z_{k-1} + \eta \nabla \log \gamma_k(z_{k-1}), 2\eta \text{Id}). \quad (10)$$

Note that sampling q_ϕ^K boils down to sampling $Z_0 \sim q_\phi$ and defining the Markov chain $\{Z_k\}_{k \leq K}$ recursively by

$$Z_k = Z_{k-1} + \eta \nabla \log \gamma_k(Z_{k-1}) + \sqrt{2\eta}U_k, \quad (11)$$

where $U_k \sim \mathcal{N}(0, \text{Id}_d)$. Moreover, as the continuous Langevin dynamics is reversible w.r.t. γ_k , this suggests that we can approximate the reversal ℓ_{k-1} of m_k by m_k directly as in AIS and thus select for any k , $\ell_{k-1} = m_k$ as done in (Heng et al., 2020). However, in this case, the weights $w_k(z_{k-1}, z_k)$ do not simplify to $\gamma_k(z_{k-1})/\gamma_{k-1}(z_{k-1})$ as m_k is not exactly γ_k -invariant and we need to rely on the general expression given in (8). This approach was concurrently and independently proposed in (Wu et al., 2020) but they do not discuss gradient computations therein.

Based on (1) and (5), we introduce

$$\mathcal{L}_{\text{SIS}} = \int \log \left(\prod_{k=1}^K w_k(z_{k-1}, z_k) \right) q_\phi^K(z_{0:K}|x) dz_{0:K}. \quad (12)$$

Algorithm 1 Langevin Monte Carlo VAE

Input: Number of steps K , initial distribution q_ϕ , unnormalized target distribution p_θ , step-size η , annealing schedule $\{\beta_k\}_{k=0}^K$.

Output: SIS estimator W of $\log p_\theta(x)$.

Draw $z_0 \sim q_\phi(\cdot|x)$;

Set $W = -\log q_\phi(z_0|x)$;

for $k = 1$ **to** K **do**

Draw $u_k \sim \varphi_d$;

Set $\gamma_k(\cdot) = \beta_k \log p_\theta(x, \cdot) + (1 - \beta_k) \log q_\phi(\cdot|x)$;

Set $z_k = z_{k-1} + \eta \nabla \log \gamma_k(z_k) + \sqrt{2\eta}u_k$;

Set $W = W + \log m_k(z_k, z_{k-1}) - \log m_k(z_{k-1}, z_k)$;

end for

Set $W = W + \log p_\theta(x, z_K)$;

Return W

We consider a reparameterization of (12) based on the Langevin mappings associated with target γ_k given by

$$\mathbb{T}_{k,u}(z_{k-1}) = z_{k-1} + \eta \nabla \log \gamma_k(z_{k-1}) + \sqrt{2\eta}u. \quad (13)$$

An easy change of variable based on the identity $Z_k = \mathbb{T}_{k,U_k}(Z_{k-1})$ in (11) shows that \mathcal{L}_{SIS} can be written as

$$\int \log \left(\prod_{k=1}^K w_k(z_{k-1}, z_k) \right) q_\phi(z_0|x) \varphi_{d,K}(u_{1:K}) dz_0 du_{1:K},$$

where φ_d stands for the density of the d -dimensional standard Gaussian distribution, $\varphi_{d,K}(u_{1:K}) = \prod_{i=1}^K \varphi_d(u_i)$, and we write $z_k = \mathbb{T}_{k,u_k} \circ \dots \circ \mathbb{T}_{1,u_1}(z_0) = \bigcirc_{i=1}^k \mathbb{T}_{i,u_i}(z_0)$. By (5) and as $\ell_{k-1} = m_k$, our objective thus finally writes

$$\begin{aligned} \mathcal{L}_{\text{SIS}}(\theta, \phi; x) &= \int q_\phi(z_0|x) \varphi_{d,K}(u_{1:K}) \\ &\times \log \left(\frac{p_\theta(x, z_K) \prod_{k=1}^K m_k(z_k, z_{k-1})}{q_\phi(z_0|x) \prod_{k=1}^K m_k(z_{k-1}, z_k)} \right) dz_0 du_{1:K}. \end{aligned} \quad (14)$$

This defines the reparameterizable Langevin Monte Carlo VAE (L-MCVAE). The algorithm to obtain an unbiased SIS estimate of $p_\theta(x)$ is described in Algorithm 1. This estimate is related to the one presented in (Caterini et al., 2018), however this work builds on a deterministic dynamics which limits the expressivity of the induced variational approximation. In contrast, we rely here on a stochastic dynamics. While we limit ourselves here to undadjusted (overdamped) Langevin dynamics, this could also be easily extended to unadjusted underdamped Langevin dynamics (Monmarché, 2020).

3. Variational Inference via Annealed Importance Sampling

In Section 2.3, we derived the SIS estimate of the evidence computed using ULA kernels M_k , whose invariant distri-

bution approximates γ_k . We can differentiate the resulting ELBO and exploit the reparametrization trick as these kernels admit a density m_k w.r.t. Lebesgue measure. When computing the AIS estimates, we need Markov kernels M_k that are γ_k -invariant. Such Markov kernels M_k most often rely on a Metropolis–Hastings (MH) step and therefore do not typically admit a transition density. While this does not invalidate the expression of the AIS estimate presented earlier, it complicates significantly the computation of an unbiased gradient estimate of the corresponding ELBO. In this section, we propose a way to compute an unbiased estimator of the ELBO gradient for MH Markov kernels. We use here elementary measure-theoretical notations which are recalled in Appendix A

3.1. Differentiating Markov kernels

Markov kernel. Let $\mathcal{B}(\mathbb{R}^d)$ denote the Borelian σ -field associated to \mathbb{R}^d . A Markov kernel M is a function defined on $\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)$, such that for any $z \in \mathbb{R}^d$, $M(z, \cdot)$ is a probability distribution on $\mathcal{B}(\mathbb{R}^d)$, *i.e.* $M(z, A)$ is the probability starting from z to hit the set $A \subset \mathbb{R}^d$. The simplest is “deterministic”, in which case $Q(z, A) = \delta_{T(z)}(A)$, where T is a measurable mapping on \mathbb{R}^d and δ_y is the Dirac mass at y . Instead of a single mapping T , we can consider a family of “indexed” mappings $\{T_u : u \in \mathbb{R}^{d_u}\}$. If g is a p.d.f on \mathbb{R}^{d_u} , we consider

$$M(z, A) = \int \mathbb{1}_A(T_u(z))g(u)du .$$

To sample $M(z, \cdot)$, we first sample $u \sim g$ and then apply the mapping $T_u(z)$. If $d = d_u$, we consider the function $G_z : \mathbb{R}^d \mapsto \mathbb{R}^d$ defined for all $z \in \mathbb{R}^d$ by

$$G_z : u \mapsto G_z(u) = T_u(z) \quad (15)$$

If G_z is a diffeomorphism, then by applying the change of variables formula, we obtain

$$\begin{aligned} M(z, A) &= \int \mathbb{1}_A(G_z(u))g(u)du \quad (16) \\ &= \int \mathbb{1}_A(z')m(z, z')dz' , \end{aligned}$$

where, denoting $J_{G_z^{-1}}(z')$ is the absolute value of the Jacobian determinant of G_z^{-1} evaluated at z' , we set

$$m(z, z') = J_{G_z^{-1}}(z')g(G_z^{-1}(z')) . \quad (17)$$

In this case, the Markov kernel has a transition density $m(z, z')$. This is the setting considered in the previous section.

Finally, some Markov kernels have both a deterministic and a continuous component. This is for example the case of Metropolis–Hastings (MH) kernels:

$$\begin{aligned} M(z, A) &= \int_{\mathcal{U}} Q_u(z, A)g(u)du , \quad \text{where} \quad (18) \\ Q_u(z, A) &= \alpha_u(z)\delta_{T_u(z)}(A) + \{1 - \alpha_u(z)\}\delta_z(A) , \end{aligned}$$

with $\alpha_u(z) = \alpha(z, T_u(z))$ is the *acceptance function* and $T_u : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the *proposal mapping*. In the sequel, we denote $\alpha_u^1(z) = \alpha_u(z)$ and $\alpha_u^0(z) = 1 - \alpha_u(z)$, and set $T_u^0(z) = z$. With these notations, (18) can be rewritten in a more concise way

$$Q_u(z, A) = \sum_{\alpha=0}^1 \alpha_u^\alpha(z)\delta_{T_u^\alpha(z)}(A) . \quad (19)$$

To sample $M(z, \cdot)$, we first draw $u \sim g$ and then compute the proposal $y = T_u(z)$. With probability $\alpha_u(z)$, the next state of the chain is set to $z' = y$, and $z' = z$ otherwise. If G_z defined in (15) is a diffeomorphism, then the Metropolis–Hasting kernel may be expressed as

$$\begin{aligned} M(z, A) &= \int \alpha(z, z')m(z, z')\mathbb{1}_A(z')dz' \\ &\quad + (1 - \bar{\alpha}(z))\delta_z(A) , \end{aligned}$$

where $\bar{\alpha}(z) = \int \alpha(z, z')m(z, z')dz'$ is the mean acceptance probability at z (the probability of accepting a move) and $m(z, z')$ is defined as in (17). In MH-algorithms, the acceptance function $\alpha : \mathbb{R}^{2d} \rightarrow [0, 1]$ is chosen so that M is π -reversible $\pi(dz)M(z, dz') = \pi(dz')M(z', dz)$, where π is the target distribution. This implies, in particular, that M is π -invariant. Standard MH algorithms use $\alpha(z, z') = 1 \wedge \pi(z')m(z', z)/\pi(z)m(z, z')$; see (Tierney, 1994).

To illustrate these definitions and constructions, consider first the symmetric Random Walk Metropolis Algorithm (RWM). In this case, $d_U = d$ and $g \leftarrow \varphi_d$, where φ_d is the d -dimensional standard Gaussian density. The proposal mapping is given by

$$G_z^{\text{RWM}}(u) = T_u^{\text{RWM}}(z) = z + \Sigma^{1/2}u ,$$

where Σ is a positive definite matrix, and the acceptance function is given by $\alpha_u^{\text{RWM}}(z) = 1 \wedge (\pi(T_u^{\text{RWM}}(z))/\pi(z))$.

Consider now the Metropolis Adjusted Langevin Algorithm (MALA); see (Besag, 1994). Assume that $z \mapsto \log \pi(z)$ is differentiable and denote by $\nabla \log \pi(z)$ its gradient. The Langevin proposal mapping T_u^{MALA} is defined by

$$G_z^{\text{MALA}}(u) = T_u^{\text{MALA}}(z) = z + \eta \nabla \log \pi(z) + \sqrt{2\eta}u . \quad (20)$$

We set $g \leftarrow \varphi_d$ and $\alpha_u^{\text{MALA}}(z)$ is the acceptance given by

$$\alpha_u^{\text{MALA}}(z) = 1 \wedge \frac{\pi(T_u^{\text{MALA}}(z))m_\eta(T_u^{\text{MALA}}(z), z)}{\pi(z)m_\eta(z, T_u^{\text{MALA}}(z))} , \quad (21)$$

where $m_\eta(z, z') = \eta^{-1/2}\varphi_d(\eta^{-1/2}\{z' - T_0^{\text{MALA}}(z)\})$, similarly to (10).

Lemma 1. For all $z \in \mathbb{R}^d$, G_z^{MALA} is a C^1 -diffeomorphism. Moreover, assume that $\log \pi$ is L -smooth with $L > 0$, *i.e.* for

$z, z' \in \mathbb{R}^d$, $\|\nabla \log \pi(z') - \nabla \log \pi(z)\| \leq L\|z' - z\|$. Then, if $0 \leq \eta < L^{-1}$, for all $u \in \mathbb{R}^d$, $\mathbb{T}_u^{\text{MALA}}$ defined in (20) is a C^1 -diffeomorphism.

The proof of Lemma 1 is postponed to Appendix C.

3.2. Differentiable AIS-based ELBO

We now generalize the derivation of Section 2.1 to handle the case where M_k and its reversal do not admit transition densities. In this case, the proposal and unnormalized target distributions are defined by

$$\begin{aligned} Q_\phi^K(dz_{0:K}|x) &= q_\phi(z_0|x) dz_0 \prod_{k=1}^K M_k(z_{k-1}, dz_k), \quad (22) \\ P_\theta^{\text{un}}(x, dz_{0:K}) &= p_\theta(x, z_K) dz_K \prod_{k=K}^1 L_{k-1}(z_k, dz_{k-1}), \end{aligned}$$

where we define the reversal Markov kernel L_{k-1} by

$$\begin{aligned} \gamma_k(z_{k-1}) dz_{k-1} M_k(z_{k-1}, dz_k) \\ = \gamma_k(z_k) dz_k L_{k-1}(z_k, dz_{k-1}). \end{aligned}$$

We consider then the AIS estimator presented in Section 2.1 in (5) by sampling independently $z_{0:K}^i \sim Q_\phi^K(\cdot|x)$ and with $w^K(z_{0:K}) = \prod_{k=1}^K w_k(z_{k-1})$, $w_k(z_{k-1}) = \gamma_k(z_{k-1})/\gamma_{k-1}(z_{k-1})$. A rigorous proof of the unbiasedness of the resulting estimator can be found in the Supplementary Material and is based on the formula

$$p_\theta(x) = \int w^K(z_{0:K}) Q_\phi^K(dz_{0:K}|x). \quad (23)$$

In the sequel, we use Metropolis Adjusted Langevin Algorithm (MALA) kernels $\{M_k\}_{k=1}^K$ targeting γ_k for each $k \in \{1, \dots, K\}$. By construction, the Markov kernel M_k is reversible w.r.t. γ_k and we set for the reversal kernel $L_{k-1} = M_k$. Note that we could easily generalize to other cases, especially inspired by recent works on non-reversible MCMC algorithms (Thin et al., 2020).

For $k \in \{1, \dots, K\}$, we use the representation of MALA kernel M_k outlined in (18) with proposal mapping $\mathbb{T}_{k,u}$ and acceptance function $\alpha_{k,u}$ defined as in (20) and (21) with $\pi \leftarrow \gamma_k$. We set

$$Q_{k,u}(z, dz') = \sum_{a=0}^1 \alpha_{k,u}^a(z) \delta_{\mathbb{T}_{k,u}^a(z)}(dz'). \quad (24)$$

By construction, the MALA kernel M_k (see (18)) writes $M_k(z, dz') = \int Q_{k,u}(z, dz') \varphi_d(u) du$. Plugging this representation into (22), we get

$$\begin{aligned} Q_\phi^K(dz_{0:K}|x) &= \int \prod_{k=1}^K Q_{k,u_k}(z_{k-1}, dz_k) \\ &\quad \times q_\phi(z_0|x) \varphi_{d,K}(u_{1:K}) dz_0 du_{1:K}, \end{aligned}$$

writing $\varphi_{d,K}(u_{1:K}) = \prod_{i=1}^K \varphi_d(u_i)$. Eq. (24) suggests to consider the extended distribution on $(z_{0:K}, a_{1:K}, u_{1:K})$:

$$\begin{aligned} Q_\phi^K(dz_{0:K}, a_{1:K}, du_{1:K}|x) &= q_\phi(z_0|x) \prod_{k=1}^K \alpha_{k,u_k}^{a_k}(z_{k-1}) \\ &\quad \times \prod_{k=1}^K \delta_{\mathbb{T}_{k,u_k}^{a_k}(z_{k-1})}(dz_k) \varphi_{d,K}(u_{1:K}) dz_0 du_{1:K}, \end{aligned}$$

which admits again as a marginal $Q_\phi^K(dz_{0:K}|x)$. Note that the variables $a_{1:K}$ correspond to the binary outcomes of the K A/R steps. By construction, $z_{1:K}$ are *deterministic functions* of $z_0, a_{1:K}$ and $u_{1:K}$: for each $k \in \{1, \dots, K\}$, $z_k = \bigcirc_{i=1}^k \mathbb{T}_{i,u_i}^{a_i}(z_0)$.

Set $w^K(z_{0:K}) = \prod_{k=1}^K w_k(z_{k-1})$, and

$$\begin{aligned} A(z_0, a_{1:K}, u_{1:K}) &= \prod_{k=1}^K \alpha_{k,u_k}^{a_k} \left(\bigcirc_{i=1}^{k-1} \mathbb{T}_{i,u_i}^{a_i}(z_0) \right), \\ W(z_0, a_{1:K}, u_{1:K}) &= \log \prod_{k=1}^K w_k \left(\bigcirc_{i=1}^{k-1} \mathbb{T}_{i,u_i}^{a_i}(z_0) \right). \end{aligned}$$

Given z_0 and $u_{1:K}$, $A(z_0, a_{1:K}, u_{1:K})$ is the conditional distribution of the A/R random variables $a_{1:K}$. It is easy to sample this distribution recursively: for each $k \in \{1, \dots, K\}$ we sample a_k from a Bernoulli distribution of parameter $\alpha_{k,u_k}(z_{k-1})$ and we set $z_k = \mathbb{T}_{k,u_k}^{a_k}(z_{k-1})$. Eqs. (5) and (9) naturally lead us to consider the ELBO

$$\begin{aligned} \mathcal{L}_{\text{AIS}} &= \int \sum_{a_{1:K}} Q_\phi^K(dz_{0:K}, a_{1:K}, du_{1:K}|x) \log [w^K(z_{0:K})] \\ &= \int \sum_{a_{1:K}} q_\phi(z_0|x) A(z_0, a_{1:K}, u_{1:K}) \\ &\quad \times W(z_0, a_{1:K}, u_{1:K}) \varphi_{d,K}(u_{1:K}) dz_0 du_{1:K}, \end{aligned}$$

where in the last line we have integrated w.r.t. $z_{1:K}$. The fact that \mathcal{L}_{AIS} is an ELBO stems immediately from (23), by applying Jensen's inequality. We can optimize this ELBO w.r.t. the different parameters at stake, even possibly the parameters of the proposal mappings.

We reparameterize the latent variable distribution $q_\phi(z_0|x)$ in terms of a known base distribution and a differentiable transformation (such as a location-scale transformation). Assuming for simplicity that $q_\phi(z_0|x)$ is a Gaussian distribution $\mathcal{N}(z_0; \mu_\phi(x), \sigma_\phi^2(x))$, then the location-scale transformation using the standard Normal as a base distribution allows us to reparameterize z_0 as $z_0 = \mu_\phi(x) + \sigma_\phi(x) \cdot u_0 = V_{\phi,x}(u_0)$, with $u_0 \sim \varphi_d$. Using this reparameterization

trick, we can write $\nabla \mathcal{L}_{\text{AIS}} = \nabla \mathcal{L}_{\text{AIS}_1} + \nabla \mathcal{L}_{\text{AIS}_2}$ with

$$\begin{aligned} \nabla \mathcal{L}_{\text{AIS}_1} &= \int \sum_{a_{1:K}} A(V_{\phi,x}(u_0), a_{1:K}, u_{1:K}) \\ &\times \nabla W(V_{\phi,x}(u_0), a_{1:K}, u_{1:K}) \varphi_{d,K+1}(u_{0:K}) du_{0:K}, \\ \nabla \mathcal{L}_{\text{AIS}_2} &= \int \sum_{a_{1:K}} A(V_{\phi,x}(u_0), a_{1:K}, u_{1:K}) W(V_{\phi,x}(u_0), a_{1:K}, u_{1:K}) \\ &\times \nabla \log A(V_{\phi,x}(u_0), a_{1:K}, u_{1:K}) \varphi_{d,K+1}(u_{0:K}) du_{0:K}. \end{aligned} \quad (25)$$

The estimation of $\nabla \mathcal{L}_{\text{AIS}_1}$ is straightforward. We sample n independent samples $u_{0:K}^{1:n} \sim \varphi_{d,K+1}$ and, for $i \in \{1, \dots, n\}$, we set $z_0^i = V_{\phi,x}(u_0^i)$ and then, for $k \in \{1, \dots, K\}$, we sample the A/R variable $a_k^i \sim \text{Ber}\{\alpha_{k,u_k^i}^1(z_{k-1}^i)\}$ and set $z_k^i = T_{k,u_k^i}^{a_k^i}(z_{k-1}^i)$, see Algorithm 2. Similarly, we then compute

$$\widehat{\nabla \mathcal{L}_{\text{AIS}_2, n}} = n^{-1} \sum_{i=1}^n \nabla W(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i).$$

The expression $\nabla \mathcal{L}_{\text{AIS}_2}$ is the REINFORCE gradient estimator (Williams, 1992) for the A/R probabilities. Indeed, we have to compute the gradient of the conditional distribution of the A/R variables given $(z_0, u_{0:K})$, and there is no obvious reparametrization for such purpose (see however (Maddison et al., 2016) for a possible solution to the problem; this solution was not investigated in this work). To reduce the variance of the REINFORCE estimator, we rely on control variates, in the spirit of (Mnih & Rezende, 2016). For $i \in \{1, \dots, n\}$, we define

$$\tilde{W}_{n,i} = \frac{1}{n-1} \sum_{j \neq i} W(V_{\phi,x}(u_0^j), a_{1:K}^j, u_{1:K}^j),$$

which is independent of $W(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i)$ and $\nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i)$ by construction. This provides the new unbiased estimator of the gradient using

$$\begin{aligned} \widehat{\nabla \mathcal{L}_{\text{AIS}_2, n}} &= n^{-1} \sum_{i=1}^n \left[W(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) - \tilde{W}_{n,i} \right] \\ &\times \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i). \end{aligned} \quad (26)$$

Algorithm 2 shows how to compute W and $\log A$.

4. Experiments

4.1. Methods and practical guidelines

In what follows, we consider two sets of experiments¹. In the first one, we aim at illustrating the expressivity and the efficiency of our estimator for VI. In the second, we tackle the problem of learning VAE: (a) Classical VAE based

¹The code to reproduce all of the experiments is available online at <https://github.com/premolab/metflow/>.

Algorithm 2 Annealed Importance Sampling VAE

Input: Number of steps K , proposal mappings $\{T_{k,u}\}_{k \leq K, u \in \mathcal{U}}$, acceptance functions $\{\alpha_{k,u}\}_{k \leq K, u \in \mathcal{U}}$, initial distribution q_ϕ , unnormalized target distribution p_θ , annealing schedule $\{\beta_k\}_{k=0}^K$.
Output: AIS estimator W of $\log p_\theta(x)$, sum $\log A$ of the A/R log probabilities.

```

Draw  $z_0 \sim q_\phi(\cdot|x)$ ;
Set  $W = 0$ ;
Set  $\log A = 0$ ;
for  $k = 1$  to  $K$  do
    Draw  $u_k \sim \varphi_d$ ;
    Draw  $a_k \sim \text{Ber}(\alpha_{k,u_k}(z_{k-1}))$ ;
    if  $a_k = 1$  (Accept) then
        Set  $z_k = z_{k-1} + \eta \nabla \log \gamma_k(z_k) + \sqrt{2\eta} u_k$ ;
    else
         $z_k = z_{k-1}$ ;
    end if
    Compute  $\log w_k(z_{k-1}) =$ 
         $(\beta_k - \beta_{k-1})(\log p_\theta(x, z_{k-1}) - \log q_\phi(z_{k-1}|x))$ ;
    Set  $W = W + \log w_k(z_{k-1})$ ;
    Set  $\log A = \log A + \log \alpha_{k,u_k}^{a_k}(z_{k-1})$ ;
end for
Return  $W, \log A$ 
    
```

on mean-field approximation (Kingma & Welling, 2013); (b) Importance-weighted Autoencoder (IWAE, (Burda et al., 2015)); (c) L-MCVAE given by Algorithm 1; (d) A-MCVAE given by Algorithm 2. We provide in the following some guidelines on how to tune the step sizes and the annealing schedules in Algorithm 1 and Algorithm 2.

A crucial hyperparameter of our method is the step size η . In principle, it could be learned by including it as an additional inference parameter ϕ and by maximizing the ELBO. However, it is then difficult to find a good trade-off between having a high A/R ratio and a large step size η at the same time. Instead, we suggest adjusting η by targeting a fixed A/R ratio ρ . It has proven effective to use a preconditioned version of (11), *i.e.* $Z_k = Z_{k-1} + \eta \odot \nabla \log \gamma_k(Z_{k-1}) + \sqrt{2\eta} \odot U_k$ with $\eta \in \mathbb{R}^p$, where we adapt each component of η using the following rule

$$\eta^{(i)} = 0.9\eta^{(i)} + 0.1\eta_0 / (\epsilon + \text{std}[\partial_{z^{(i)}} \log p_\theta(x, z)]) . \quad (27)$$

Here std denotes the standard deviation over the batch x of the quantity $\partial_{z^{(i)}} \log p_\theta(x, z)$, and $\epsilon > 0$. The scalar η_0 is a tuning parameter which is adjusted to target the A/R ratio ρ . This strategy follows the same heuristics as Adam (Kingma & Ba, 2014). In the following ρ is set to 0.8 for A-MCVAE and 0.9 for L-MCVAE (keeping it high for L-MCVAE ensures that the Langevin dynamics stays “almost reversible”, thus keeping a low variance SIS estimator).




Figure 1. Visualization of the posterior approximation given after optimization of different bounds for toy generation process. Top row, from left to right: True posterior, VAE posterior, IWAE posterior. Bottom row, from left to right: VI with RealNVP posterior, A-MCVAE posterior, L-MCVAE posterior.

An optimal choice of the temperature schedule $\{\beta_k\}_{k=0}^K$ for SIS and AIS is a difficult problem. We have focused in our experiments on three different settings. First, we consider the temperature schedule fixed and regularly spaced between 0 and 1. Following (Grosse et al., 2015), the second option is the sigmoidal tempering scheme where $\beta_k = (\tilde{\beta}_k - \tilde{\beta}_1)/(\tilde{\beta}_K - \tilde{\beta}_1)$ with, $\tilde{\beta}_k = \sigma(\delta(2k/K - 1))$, σ is the sigmoid function and $\delta > 0$ is a parameter that we optimize during the training phase. The last schedule consists in learning the temperatures $\{\beta_k\}_{k=0}^K$ directly as additional inference parameters ϕ .

4.2. Toy 2D example and Probabilistic Principal Component Analysis

In the following two examples, we fix the parameters θ of the likelihood model and apply Algorithm 1 and Algorithm 2 to perform VI to sample from $z \mapsto p_\theta(z|x)$. Consider first a toy hierarchical example where we generate some i.i.d. data $x = (x_i)_{i=1}^N \in \mathbb{R}^N$ from the i.i.d. latent variables $z = (z_i)_{i=1}^N \in \mathbb{R}^{2N}$ as follows for $\xi > 0$ $x_i|z_i \sim \mathcal{N}(\xi \cdot (\|z_i\|^2 + \zeta), \sigma^2) = p_\theta(x_i|z_i)$. We consider the variational approximation as $q_\phi(z|x) = \mathcal{N}(z; \mu_\phi(x), \sigma_\phi(x)^2 \text{Id})$, where $\mu_\phi(x), \sigma_\phi(x) \in \mathbb{R}^d$ are the outputs of a fully connected neural network with parameters ϕ . We compare these algorithms to VI using Real-valued Non-Volume Preserving transformation (RealNVP, Dinh et al. (2016)).

Figure 1 displays the VI posterior approximations corresponding to the different schemes for a given observation x . It can be observed that MCVAE benefit from flexible variational distributions compared to other classical schemes, which mostly fail to recover the true posterior. Additional results on the estimation of the parameters ξ, ζ , given in the supplementary material, further support our claims; see Appendix B.1.

We now illustrate the performance of MCVAE on a probabilistic principal component analysis problem applied to MNIST (Salakhutdinov & Murray, 2008), as we can access in this case the exact likelihood and its gradient. We follow the set-up of (Ruiz et al., 2021, Section 6.1). We consider here a batch of size $N = 100$ for the




Figure 2. Representation of the different estimators (top) and their gradient (bottom) of the true log likelihood. From left to right, a/ L-MCVAE, $K = 5$, b/ L-MCVAE, $K = 10$, c/ A-MCVAE, $K = 5$, d/ A-MCVAE, $K = 10$, e/ A-MCVAE, $K = 5$ with control variates.

model $p_\theta(x, z) = \mathcal{N}(z; 0, \text{Id}_d)\mathcal{N}(x; \theta_0 + \theta_1 z, \sigma^2 \text{Id}_p)$, with $d = 100$ and $p = 784$. We fix arbitrarily θ_0, θ_1 , and fit an amortized variational distribution $q_\phi(z|x)$ by maximizing the IWAE bound w.r.t. ϕ with $K = 100$ importance samples for a large number of epochs. The distribution $q_\phi(z|x) = \mathcal{N}(z; \mu_\phi(x), \text{diag}(\sigma_\phi^2(x)))$ is a mean-field Gaussian distribution where μ_ϕ, σ_ϕ are linear functions of the observation x .

We compare the Langevin SIS estimator (L-MCVAE) of the log evidence $\log p_\theta(x)$ with Langevin auxiliary kernels as described in Section 2.3, and the Langevin AIS estimate (A-MCVAE). Moreover, we also compare the gradients of these quantities w.r.t. the parameters θ_0, θ_1 .

Figure 2 summarises the results with boxplots computed over 200 independent samples of each estimator. The quantity reported on the first boxplot corresponds to the Monte Carlo samples of $\log \hat{p}_\theta(x) - \log p_\theta(x)$. On the one hand, we note that the SIS estimator has larger variance than AIS, and that the latter achieves a better ELBO. Moreover, in both cases, increasing the number of steps K tightens the

bound. On the other hand, the estimator of the gradient of AIS is noisier than that of SIS, even though variance reduction techniques allows us to recover a similar variance. We also present in the supplementary material the Langevin SIS estimator using auxiliary backward kernels learnt with neural networks (as done in previous contributions); see Appendix B.2. The auxiliary neural backward kernels are set as $l(z, z') = \mathcal{N}(z'; \mu_\psi(z), \text{diag}(\sigma_\psi^2(z)))$, $\mu_\psi, \sigma_\psi \in \mathbb{R}^d$, where the parameters ψ are learnt through the SIS ELBO, similarly to (Huang et al., 2018). The variance of the associated estimator and their gradients are larger than that of SIS using the approximate reversals as backward kernels; i.e. $\ell_{k-1} = m_k$.




Figure 3. Log-likelihood of L-MCVAE depending on the number of Langevin steps K . Increasing K improves performance, however at the expense of the computational complexity.

4.3. Numerical results for image datasets

Following (Wu et al., 2016), we propose to evaluate our models using AIS (not to be confused with the proposed AIS-based VI approach) to get an estimation of the negative log-likelihood. The base distribution is the distribution output by the encoder, and we perform K steps of annealing to compute the estimator of the likelihood, as given by (5). In practice, we use $K = 5$ HMC steps with 3 leapfrogs for evaluating our models.




Figure 4. Evolution of the held-out loglikelihood during training for A-MCVAE, L-MCVAE, IWAE and VAE on MNIST.

We evaluate our models on three different datasets: MNIST, CIFAR-10 and CelebA. All the models we compare share the same architecture: the inference network q_ϕ is given by a convolutional network with 8 convolutional layers and one linear layer, which outputs the parameters $\mu_\phi(x), \sigma_\phi(x) \in \mathbb{R}^d$ of a factorized Gaussian distribution, while the generative model $p_\theta(\cdot|z)$ is given by another convolutional network π_θ , where we use nearest neighbor upsamplings. This outputs the parameters for the factorized Bernoulli distribution (for MNIST dataset), that is

$$p_\theta(x|z) = \prod_{i=1}^N \text{Ber}(x^{(i)} | (\pi_\theta(z))^{(i)})$$

and similarly the mean of the Gaussian distributions for colored datasets (CIFAR-10, Celeba). We compare A-MCVAE, L-MCVAE, IWAE, and VAE with different settings. All the models are implemented using PyTorch (Paszke et al., 2019) and optimized using the Adam optimizer (Kingma & Ba, 2014) for 100 epochs each. The training process is using PyTorch Lightning toolkit (Falcon, 2019).

First, consider dynamically binarized MNIST dataset (Salakhutdinov & Murray, 2008). In this case, the latent dimension is set to $d = 64$. We present in Table 1 the results of the different models at different stages of the optimization. Moreover, we show on Figure 3 the performance of L-MCVAE for different values of K compared to IWAE baseline. In particular, we see that increasing K increases the performance of our VAE, however at the expense of an increase in computational cost. We also display on Figure 4 the evolution of the held-out loglikelihood for various objectives during training. Adding Langevin transitions appears to help convergence of the models.

Second, we compare similarly the different models on CelebA and CIFAR, see Table 2 and Table 3. In this case, the latent dimension is chosen to be $d = 128$. Increasing the number of MCMC steps seems again to improve both the ELBO and the final loglikelihood estimate. In each case, all models are run with 5 different seeds to compute the presented empirical standard deviation.

5. Discussion

We have shown in this article how one can leverage state-of-the-art Monte Carlo estimators of the evidence to develop novel competitive VAEs by developing novel gradient estimates of the corresponding ELBOs.

For a given computational complexity, AIS based on MALA provides ELBO estimates which are typically tighter than SIS estimates based on ULA. However, the variance of the gradient estimates of the AIS-based ELBO (A-MCVAE) is also significantly larger than for the SIS-based ELBO

Table 1. Results of the different models on MNIST. A more detailed version of this table is included in the supplementary material.

number of epochs	negative ELBO estimate			NLL estimate		
	10	30	100	10	30	100
VAE	95.26 ± 0.49	91.58 ± 0.27	89.70 ± 0.19	89.83 ± 0.59	86.86 ± 0.26	85.22 ± 0.07
IWAE, $K = 10$	91.42 ± 0.21	88.56 ± 0.07	87.16 ± 0.19	88.54 ± 0.27	86.07 ± 0.1	84.82 ± 0.1
IWAE, $K = 50$	90.34 ± 0.27	87.5 ± 0.16	86.05 ± 0.11	89.4 ± 0.25	86.54 ± 0.15	85.05 ± 0.1
L-MCVAE, $K = 5$	96.62 ± 3.24	88.58 ± 0.75	87.51 ± 0.41	90.59 ± 2.01	85.68 ± 0.49	84.92 ± 0.24
L-MCVAE, $K = 10$	96.78 ± 1.06	87.99 ± 0.71	86.8 ± 0.66	91.33 ± 0.61	85.47 ± 0.46	84.58 ± 0.39
A-MCVAE, $K = 3$	96.21 ± 3.43	88.64 ± 0.78	87.63 ± 0.42	90.42 ± 2.34	85.77 ± 0.65	85.02 ± 0.37
A-MCVAE, $K = 5$	95.55 ± 2.96	87.99 ± 0.57	87.03 ± 0.27	90.39 ± 2.21	85.6 ± 0.67	84.84 ± 0.38
VAE with RealNVP	95.23 ± 0.33	91.69 ± 0.15	89.62 ± 0.17	89.98 ± 0.24	86.88 ± 0.05	85.23 ± 0.18

Table 2. Results of the different models on CelebA. A more detailed version of this table is included in the supplementary material. 11400 must be added to all scores in this table.

number of epochs	negative ELBO - 11400+			NLL - 11400+		
	10	30	100	10	30	100
VAE	23.78 ± 1.95	17.99 ± 0.4	14.72 ± 0.16	17.35 ± 1.7	12.68 ± 0.62	10.11 ± 0.32
IWAE, $K = 10$	20.59 ± 0.71	15.45 ± 0.52	12.2 ± 0.3	18.25 ± 0.6	13.18 ± 0.42	10.14 ± 0.31
IWAE, $K = 50$	19.05 ± 0.39	13.59 ± 0.5	10.48 ± 0.89	19.08 ± 0.42	13.17 ± 0.54	10.12 ± 0.86
L-MCVAE, $K = 5$	21.61 ± 1.48	12.72 ± 0.43	11.6 ± 0.37	16.42 ± 1.47	9.62 ± 0.47	8.72 ± 0.4
L-MCVAE, $K = 10$	20.7 ± 1.15	11.81 ± 0.34	10.6 ± 0.23	17.0 ± 1.87	9.29 ± 0.73	8.24 ± 0.52
A-MCVAE, $K = 3$	21.59 ± 1.5	13.94 ± 0.42	12.84 ± 0.3	16.64 ± 1.37	10.98 ± 0.48	9.95 ± 0.3
A-MCVAE, $K = 5$	20.95 ± 1.18	12.42 ± 0.42	11.13 ± 0.37	17.42 ± 1.49	9.97 ± 0.59	8.82 ± 0.57
VAE with RealNVP	15.12 ± 0.48	13.63 ± 0.27	12.58 ± 0.61	10.42 ± 0.33	9.04 ± 0.26	8.98 ± 0.2

Table 3. Results of the different models on CIFAR. A more detailed version of this table is included in the supplementary material. 2800 must be added to all scores in this table.

number of epochs	negative ELBO - 2800+			NLL - 2800+		
	10	30	100	10	30	100
VAE	69.57 ± 0.08	69.55 ± 0.51	68.84 ± 0.06	68.51 ± 0.07	68.41 ± 0.33	67.9 ± 0.03
IWAE, $K= 10$	69.82 ± 0.03	69.35 ± 0.03	69.36 ± 0.36	68.56 ± 0.03	68.0 ± 0.03	68.02 ± 0.4
IWAE, $K= 50$	69.94 ± 0.08	69.55 ± 0.04	69.43 ± 0.03	69.15 ± 0.15	68.37 ± 0.18	67.93 ± 0.02
L-MCVAE, $K= 5$	70.62 ± 0.41	68.55 ± 0.18	68.09 ± 0.1	69.15 ± 0.38	67.73 ± 0.07	67.5 ± 0.07
L-MCVAE, $K= 10$	70.99 ± 0.59	68.36 ± 0.04	68.03 ± 0.0	69.8 ± 0.67	67.76 ± 0.04	67.51 ± 0.03
A-MCVAE, $K= 3$	69.97 ± 0.99	68.48 ± 0.29	68.18 ± 0.16	69.26 ± 0.76	67.77 ± 0.18	67.55 ± 0.1
A-MCVAE, $K= 5$	70.1 ± 0.89	68.28 ± 0.2	68.01 ± 0.08	69.23 ± 0.75	67.71 ± 0.15	67.5 ± 0.07
VAE with RealNVP	70.01 ± 0.12	69.51 ± 0.07	69.19 ± 0.13	68.73 ± 0.05	68.35 ± 0.05	68.05 ± 0.02

(L-MCVAE) as it has to rely on REINFORCE gradient estimates. While control variates can be considered to reduce the variance, this comes at a significant increase in computational cost.

Empirically, L-MCVAE should thus be favoured as it provides both a tighter ELBO than standard techniques and low variance gradient estimates.

Acknowledgements

The work was partly supported by ANR-19-CHIA-0002-01 ‘‘SCAI’’ and EPSRC CoSInES grant EP/R034710/1. It was partly carried out under the framework of HSE University

Basic Research Program. The development of a software system for the experimental study of VAEs and its application to computer vision problems (Section 4) was supported by the Russian Science Foundation grant 20-71-10135. Part of this research has been carried out under the auspice of the Lagrange Center for Mathematics and Computing.

References

- Besag, J. Comments on “Representations of knowledge in complex systems” by u. Grenander and m. Miller. *J. Roy. Statist. Soc. Ser. B*, 56:591–592, 1994.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. Hamiltonian variational auto-encoder. In *Advances in Neural Information Processing Systems*, pp. 8167–8177, 2018.
- Crooks, G. E. Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *Journal of Statistical Physics*, 90(5-6):1481–1487, 1998.
- Del Moral, P., Doucet, A., and Jasra, A. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B*, 68(3):411–436, 2006.
- Ding, X. and Freedman, D. J. Learning deep generative models with annealed importance sampling. *arXiv preprint arXiv:1906.04904*, 2019.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Domke, J. and Sheldon, D. R. Importance weighting and variational inference. In *Advances in neural information processing systems*, pp. 4470–4479, 2018.
- Falcon, W. Pytorch lightning. *GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>*, 3, 2019.
- Goyal, A. G. A. P., Ke, N. R., Ganguli, S., and Bengio, Y. Variational walkback: Learning a transition operator as a stochastic recurrent net. In *Advances in Neural Information Processing Systems*, pp. 4392–4402, 2017.
- Grosse, R. B., Ghahramani, Z., and Adams, R. P. Sandwiching the marginal likelihood using bidirectional monte carlo. *arXiv preprint arXiv:1511.02543*, 2015.
- Heng, J., Bishop, A. N., Deligiannidis, G., and Doucet, A. Controlled sequential Monte Carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020.
- Huang, C.-W., Tan, S., Lacoste, A., and Courville, A. C. Improving explorability in variational inference with annealed variational objectives. In *Advances in Neural Information Processing Systems*, pp. 9701–9711, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improving variational inference with inverse autoregressive flow. *arXiv preprint arXiv:1606.04934*, 2016.
- Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. Auxiliary deep generative models. In *International conference on machine learning*, pp. 1445–1453. PMLR, 2016.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Maddison, C. J., Lawson, D., Tucker, G., Heess, N., Norouzi, M., Mnih, A., Doucet, A., and Teh, Y. W. Filtering variational objectives. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6576–6586, 2017.
- Mnih, A. and Rezende, D. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, pp. 2188–2196. PMLR, 2016.
- Monmarché, P. High-dimensional MCMC with a standard splitting scheme for the underdamped Langevin. *arXiv preprint arXiv:2007.05455*, 2020.
- Neal, R. M. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- Ranganath, R., Tran, D., and Blei, D. Hierarchical variational models. In *International Conference on Machine Learning*, pp. 324–333. PMLR, 2016.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538. PMLR, 2015.
- Ruiz, F. J., Titsias, M. K., Cemgil, T., and Doucet, A. Unbiased gradient estimation for variational auto-encoders using coupled Markov chains. In *Uncertainty in Artificial Intelligence*, 2021.
- Salakhutdinov, R. and Murray, I. On the quantitative analysis of deep belief networks. In *Proceedings of the 25th international conference on Machine learning*, pp. 872–879, 2008.

-
- Salimans, T., Kingma, D., and Welling, M. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pp. 1218–1226, 2015.
- Thin, A., Kotelevskii, N., Andrieu, C., Durmus, A., Moulines, E., and Panov, M. Nonreversible MCMC from conditional invertible transforms: a complete recipe with convergence guarantees. *arXiv preprint arXiv:2012.15550*, 2020.
- Tierney, L. Markov Chains for Exploring Posterior Distributions. *The Annals of Statistics*, 22(4):1701–1728, 1994.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.
- Wu, H., Köhler, J., and Noé, F. Stochastic normalizing flows. *Advances in Neural Information Processing Systems*, 2020.
- Wu, Y., Burda, Y., Salakhutdinov, R., and Grosse, R. On the quantitative analysis of decoder-based generative models. *arXiv preprint arXiv:1611.04273*, 2016.

Monte Carlo VAE

SUPPLEMENTARY DOCUMENT

A. Notations and definitions

Let (X, \mathcal{X}) be a measurable space. A *Markov kernel* N on $X \times \mathcal{X}$ is a mapping $N : X \times \mathcal{X} \rightarrow [0, 1]$ satisfying the following conditions:

- (i) for every $x \in X$, the mapping $N(x, \cdot) : A \mapsto N(x, A)$ is a probability of on \mathcal{X} ,
- (ii) for every $A \in \mathcal{X}$, the mapping $N(\cdot, A) : x \mapsto N(x, A)$ is a measurable function from (X, \mathcal{X}) to $([0, 1], \mathcal{B}([0, 1]))$, where $\mathcal{B}([0, 1])$ denotes the borelian sets of $[0, 1]$.

Let λ be a positive σ -finite measure on (X, \mathcal{X}) and $n : X \times X \rightarrow \mathbb{R}_+$ be a nonnegative function, measurable with respect to the product σ -field $\mathcal{X} \otimes \mathcal{X}$. Then, the application N defined on $X \times \mathcal{X}$ by

$$N(x, A) = \int_A n(x, y) \lambda(dy),$$

is a kernel. The function n is called the density of the kernel N w.r.t. the measure λ . The kernel N is Markovian if and only if $\int_X n(x, y) \lambda(dy) = 1$ for all $x \in X$.

Let N be a kernel on $X \times \mathcal{X}$ and f be a nonnegative function. A function $Nf : X \rightarrow \mathbb{R}_+$ is defined by setting, for $x \in X$,

$$Nf(x) = \int_X N(x, dy) f(y).$$

Let μ be a probability on (X, \mathcal{X}) . For $A \in \mathcal{X}$, define

$$\mu N(A) = \int_X \mu(dx) N(x, A).$$

If N is Markovian, then μN is a probability on (X, \mathcal{X}) .

B. Experiences

B.1. Toy example

We first describe additional experiments on the toy dataset introduced in Section 4.2.

Recall that we generate some i.i.d. data $x = (x_i)_{i=1}^N \in \mathbb{R}^N$ from the i.i.d. latent variables $z = (z_i)_{i=1}^N \in \mathbb{R}^{2N}$ as follows for $\eta > 0$: $z_i \sim N(0; \text{Id})$ and $x_i | z_i \sim N(\eta \cdot (\|z_i\| + \zeta), \sigma^2) = p_\theta(x_i | z_i)$.

This example, presented for $z \in \mathbb{R}^2$, easily extends to the case where z lies in \mathbb{R}^d , with d increasing from 2 to 300. We tackle here the problem at estimating the parameter $\theta = (\eta, \zeta)$ when d varies.

We show in Figure S1 the error $\|\hat{\theta} - \theta\|^2$ for the different methods. The increased flexibility of the posterior proves more effective for estimating the true parameters of the generative model.

B.2. Probabilistic Principal Component Analysis

We detail the impact of the learnable reverse kernels on the variance of the estimator and looseness of the ELBO. In our experiments, reverse kernels were given by fully-connected neural networks. We train K different reverse kernels $\{l_k\}_{k=0}^{K-1}$




Figure S1. Squared error for parameter’s estimates, obtained using different models.




Figure S2. Representation of the different estimators (left) and their gradient (right) of the true log likelihood. From left to right, a/ L-MCVAE, $K = 5$, b/ L-MCVAE, $K = 10$, c/ L-MCVAE, $K = 1$, learnable reverse, d/ L-MCVAE, $K = 2$ learnable reverse, e/ A-MCVAE, $K = 5$, f/ A-MCVAE, $K = 10$, g/ A-MCVAE, $K = 5$ with control variates.

for the K transitions, each given by a separate neural network, and amortized over the observation x , similarly to (Salimans et al., 2015; Huang et al., 2018). Given the parameters (θ, ϕ) , we train these kernels for a large number of epochs using the SIS objective (14) and the Adam optimizer (Kingma & Ba, 2014). In particular, we display in Figure S2 the different estimators to be compared. It is easily seen that reverse kernels can not provide reasonable and stable density estimates. At the same time, we observe the variance of the gradient is higher in those models than in the ones we present in the main text. This motivates our approach bypassing the optimization of the reverse kernels.

B.3. Additional experimental results

We display in this section the full results on MNIST, CelebA and CIFAR respectively of the different models as well as the effect of the different annealing schemes (respectively in Table 1, Table 2 and 3).

C. Proofs

C.1. Proof of SIS and AIS Identities

Proposition S2. Let $\{\Gamma_k\}_{k=0}^K$ be a sequence of distributions on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, $\{M_k\}_{k=1}^K$ and $\{L_k\}_{k=0}^{K-1}$ be Markov kernels. Assume that for each $k \in \{0, \dots, K-1\}$, there exists a positive measurable function $w_k : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}_+$ such that

$$\Gamma_k(dz_k)L_{k-1}(z_k, dz_{k-1}) = \Gamma_{k-1}(dz_{k-1})M_k(z_{k-1}, dz_k)w_k(z_{k-1}, z_k). \quad (\text{S1})$$

Table 1. Results of the different models on MNIST with different annealing schemes.

number of epoches	<i>ELBO</i> : 10	30	100	<i>NLL</i> : 10	30	100
VAE	95.26 ± 0.5	91.58 ± 0.27	89.7 ± 0.19	89.83 ± 0.59	86.86 ± 0.26	85.22 ± 0.07
IWAE, K= 10	91.42 ± 0.21	88.56 ± 0.07	87.17 ± 0.19	88.54 ± 0.27	86.07 ± 0.1	84.82 ± 0.1
IWAE, K= 50	90.34 ± 0.27	87.5 ± 0.16	86.05 ± 0.11	89.4 ± 0.25	86.54 ± 0.15	85.05 ± 0.1
L-MCVAE Fixed, K= 5	96.6 ± 3.51	88.8 ± 0.46	87.77 ± 0.12	90.63 ± 2.19	85.85 ± 0.27	85.07 ± 0.04
L-MCVAE Sigmoidal, K= 5	95.48 ± 2.29	88.87 ± 0.82	87.81 ± 0.53	90.05 ± 1.63	85.92 ± 0.62	85.16 ± 0.38
L-MCVAE All learnable, K= 5	96.62 ± 3.24	88.58 ± 0.75	87.51 ± 0.41	90.59 ± 2.01	85.68 ± 0.49	84.92 ± 0.24
L-MCVAE Fixed, K= 10	95.98 ± 3.91	88.36 ± 0.7	87.38 ± 0.35	90.5 ± 2.23	85.75 ± 0.33	85.0 ± 0.11
L-MCVAE Sigmoidal, K= 10	96.78 ± 0.47	88.35 ± 0.63	87.17 ± 0.52	91.13 ± 0.27	85.72 ± 0.31	84.84 ± 0.26
L-MCVAE All learnable, K= 10	96.78 ± 1.06	87.99 ± 0.71	86.8 ± 0.66	91.33 ± 0.61	85.47 ± 0.46	84.58 ± 0.39
A-MCVAE Fixed, K= 3	96.21 ± 3.43	88.64 ± 0.78	87.63 ± 0.42	90.42 ± 2.34	85.77 ± 0.65	85.02 ± 0.37
A-MCVAE Sigmoidal, K= 3	96.59 ± 2.31	88.96 ± 0.4	87.86 ± 0.06	90.85 ± 1.62	85.97 ± 0.34	85.17 ± 0.1
A-MCVAE All learnable, K= 3	95.44 ± 2.68	88.79 ± 0.63	87.78 ± 0.37	89.9 ± 1.68	85.96 ± 0.59	85.23 ± 0.41
A-MCVAE Fixed, K= 5	95.55 ± 2.96	87.99 ± 0.57	87.03 ± 0.27	90.39 ± 2.21	85.6 ± 0.67	84.84 ± 0.38
A-MCVAE Sigmoidal, K= 5	96.56 ± 2.02	88.51 ± 0.31	87.46 ± 0.48	91.62 ± 1.55	85.96 ± 0.06	85.15 ± 0.21
A-MCVAE All learnable, K= 5	95.81 ± 1.72	88.11 ± 0.13	87.14 ± 0.18	90.79 ± 1.14	85.71 ± 0.28	84.95 ± 0.04
VAE with RealNVP	95.23 ± 0.33	91.69 ± 0.15	89.62 ± 0.17	89.98 ± 0.24	86.88 ± 0.05	85.23 ± 0.18

Table 2. Full results of the different models on CelebA. All scores must be added 11400 in this table.

number of epoches	<i>ELBO</i> : 10	30	100	<i>NLL</i> : 10	30	100
VAE	23.78 ± 1.95	17.99 ± 0.4	14.72 ± 0.16	17.35 ± 1.7	12.68 ± 0.62	10.11 ± 0.32
IWAE, K= 10	20.59 ± 0.71	15.45 ± 0.52	12.2 ± 0.3	18.25 ± 0.6	13.18 ± 0.42	10.14 ± 0.31
IWAE, K= 50	19.05 ± 0.39	13.59 ± 0.5	10.48 ± 0.89	19.08 ± 0.42	13.17 ± 0.54	10.12 ± 0.86
L-MCVAE Fixed, K= 5	21.93 ± 1.34	13.12 ± 1.27	12.03 ± 1.21	16.65 ± 1.55	10.12 ± 1.38	9.14 ± 1.27
L-MCVAE Sigmoidal, K= 5	21.61 ± 1.48	12.72 ± 0.43	11.6 ± 0.37	16.42 ± 1.47	9.62 ± 0.47	8.72 ± 0.4
L-MCVAE All learnable, K= 5	20.75 ± 0.65	12.99 ± 0.7	11.91 ± 0.61	16.16 ± 0.93	10.01 ± 0.72	9.03 ± 0.64
L-MCVAE Fixed, K= 10	21.49 ± 0.03	12.83 ± 0.57	11.76 ± 0.56	17.67 ± 0.75	10.26 ± 0.9	9.24 ± 0.79
L-MCVAE Sigmoidal, K= 10	19.44 ± 0.82	11.81 ± 0.45	10.7 ± 0.4	15.67 ± 1.48	9.24 ± 0.8	8.24 ± 0.73
L-MCVAE All learnable, K= 10	20.7 ± 1.15	11.81 ± 0.34	10.6 ± 0.23	17.0 ± 1.87	9.29 ± 0.73	8.26 ± 0.52
A-MCVAE Fixed, K= 3	21.59 ± 1.5	13.94 ± 0.42	12.84 ± 0.3	16.64 ± 1.37	10.98 ± 0.48	9.95 ± 0.3
A-MCVAE Sigmoidal, K= 3	23.63 ± 1.19	14.17 ± 0.26	12.96 ± 0.18	18.0 ± 0.54	11.09 ± 0.2	10.11 ± 0.13
A-MCVAE All learnable, K= 3	22.11 ± 1.66	14.62 ± 0.35	13.54 ± 0.18	17.38 ± 1.54	11.68 ± 0.33	10.67 ± 0.16
A-MCVAE Fixed, K= 5	20.13 ± 1.11	13.11 ± 0.38	11.99 ± 0.56	16.71 ± 1.47	10.64 ± 0.24	9.63 ± 0.32
A-MCVAE Sigmoidal, K= 5	20.95 ± 1.18	12.42 ± 0.42	11.13 ± 0.37	17.42 ± 1.49	9.97 ± 0.59	8.82 ± 0.57
A-MCVAE All learnable, K= 5	22.17 ± 0.17	12.73 ± 0.09	11.46 ± 0.15	18.97 ± 1.04	10.41 ± 0.28	9.22 ± 0.16
VAE with RealNVP	15.56 ± 0.29	13.60 ± 0.35	12.21 ± 0.27	10.69 ± 0.19	9.09 ± 0.26	8.98 ± 0.2

Table 3. Results of the different models on CIFAR-10 with different annealing schemes. All scores must be added 2800 in this table.

number of epoches	<i>ELBO</i> : 10	30	100	<i>NLL</i> : 10	30	100
VAE	69.57 ± 0.08	69.55 ± 0.51	68.84 ± 0.06	68.51 ± 0.07	68.41 ± 0.33	67.9 ± 0.03
IWAE, K= 10	69.82 ± 0.03	69.35 ± 0.03	69.36 ± 0.36	68.56 ± 0.03	68.0 ± 0.03	68.02 ± 0.4
IWAE, K= 50	69.94 ± 0.08	69.55 ± 0.04	69.43 ± 0.03	69.15 ± 0.15	68.37 ± 0.18	67.93 ± 0.02
L-MCVAE Fixed, K= 5	70.86 ± 0.53	68.44 ± 0.18	68.12 ± 0.11	69.37 ± 0.37	67.78 ± 0.1	67.53 ± 0.07
L-MCVAE Sigmoidal, K= 5	70.9 ± 0.59	68.46 ± 0.13	68.12 ± 0.11	69.42 ± 0.39	67.77 ± 0.11	67.51 ± 0.08
L-MCVAE All learnable, K= 5	70.62 ± 0.41	68.55 ± 0.18	68.09 ± 0.1	69.15 ± 0.38	67.73 ± 0.07	67.5 ± 0.07
L-MCVAE Fixed, K= 10	70.67 ± 0.42	68.37 ± 0.06	69.07 ± 1.49	69.62 ± 0.54	67.78 ± 0.06	67.51 ± 0.03
L-MCVAE Sigmoidal, K= 10	70.99 ± 0.59	68.36 ± 0.04	68.03 ± 0.0	69.8 ± 0.67	67.76 ± 0.04	67.51 ± 0.03
L-MCVAE All learnable, K= 10	71.19 ± 0.79	68.36 ± 0.03	68.01 ± 0.04	69.95 ± 0.62	67.78 ± 0.07	67.5 ± 0.05
A-MCVAE Fixed, K= 3	69.97 ± 0.99	68.48 ± 0.29	68.18 ± 0.16	69.26 ± 0.76	67.77 ± 0.18	67.55 ± 0.1
A-MCVAE Sigmoidal, K= 3	70.5 ± 1.18	68.45 ± 0.28	68.19 ± 0.18	69.18 ± 0.8	67.77 ± 0.19	67.56 ± 0.11
A-MCVAE All learnable, K= 3	70.69 ± 1.23	68.44 ± 0.3	68.17 ± 0.18	69.36 ± 0.89	67.76 ± 0.2	67.55 ± 0.11
A-MCVAE Fixed, K= 5	70.37 ± 1.04	68.31 ± 0.21	68.04 ± 0.1	69.36 ± 0.87	67.73 ± 0.17	67.51 ± 0.08
A-MCVAE Sigmoidal, K= 5	70.89 ± 0.38	68.4 ± 0.05	68.07 ± 0.04	69.71 ± 0.33	67.8 ± 0.04	67.53 ± 0.02
A-MCVAE All learnable, K= 5	70.1 ± 0.89	68.28 ± 0.2	68.01 ± 0.08	69.23 ± 0.75	67.71 ± 0.15	67.5 ± 0.07
VAE with RealNVP	70.01 ± 0.12	69.51 ± 0.07	69.19 ± 0.13	68.73 ± 0.05	68.35 ± 0.05	68.05 ± 0.02

Then,

$$\Gamma_0(dz_0) \prod_{k=1}^K M_k(z_{k-1}, dz_k) \prod_{k=1}^K w_k(z_{k-1}, z_k) = \Gamma_K(dz_K) \prod_{k=K}^1 L_{k-1}(z_k, dz_{k-1}). \quad (S2)$$

Proof. We prove by induction that for $k \in \{1, \dots, K\}$,

$$\Gamma_0(dz_0) \prod_{i=1}^k M_i(z_{i-1}, dz_i) \prod_{i=1}^k w_i(z_{i-1}, z_i) = \Gamma_k(dz_k) \prod_{i=k}^1 L_{i-1}(z_i, dz_{i-1}). \quad (S3)$$

Eq. (S3) is satisfied for $k = 1$ by (S1). Assume that (S3) is satisfied for $k \leq K - 1$. By (S1),

$$\begin{aligned} \Gamma_{k+1}(dz_{k+1}) \prod_{i=k+1}^1 L_{i-1}(z_i, dz_{i-1}) &= \Gamma_{k+1}(dz_{k+1}) L_k(z_{k+1}, dz_k) \prod_{i=k}^1 L_{i-1}(z_i, dz_{i-1}) \\ &= \Gamma_k(dz_k) M_{k+1}(z_k, dz_{k+1}) w_{k+1}(z_k, z_{k+1}) \prod_{i=k}^1 L_{i-1}(z_i, dz_{i-1}) \\ &= M_{k+1}(z_k, dz_{k+1}) w_{k+1}(z_k, z_{k+1}) \Gamma_0(dz_0) \prod_{i=1}^k M_i(z_{i-1}, dz_i) \prod_{i=1}^k w_i(z_{i-1}, z_i) \end{aligned}$$

which concludes the proof. \square

We now highlight conditions under which (S1) is satisfied.

1. Assume that $\{\Gamma_k\}_{k=0}^K$ have positive densities w.r.t. to the Lebesgue measure, i.e. $\Gamma_k(dz_k) = \Gamma_k(z_k) dz_k$ and that the kernels $\{M_k\}_{k=1}^K$ and $\{L_k\}_{k=0}^{K-1}$ have positive transition densities $M_k(z_{k-1}, dz_k) = m_k(z_{k-1}, z_k) dz_k$ and $L_{k-1}(z_k, dz_{k-1}) = \ell_{k-1}(z_k, z_{k-1}) dz_{k-1}$, $k \in \{1, \dots, K\}$. Then,

$$w_k(z_{k-1}, z_k) = \frac{\gamma_k(z_k) \ell_{k-1}(z_k, z_{k-1})}{\gamma_{k-1}(z_{k-1}) m_k(z_{k-1}, z_k)}$$

2. Assume that for $k \in \{1, \dots, K\}$, $\Gamma_k(dz_{k-1})M_k(z_{k-1}, dz_k) = \Gamma_k(dz_k)L_{k-1}(z_k, dz_{k-1})$, and that there exists a positive measurable function such that $\Gamma_k(dz_{k-1}) = \tilde{w}_k(z_{k-1})\Gamma_{k-1}(dz_{k-1})$. Then,

$$\Gamma_k(dz_k)L_{k-1}(z_k, dz_{k-1}) = \Gamma_k(dz_{k-1})M_k(z_{k-1}, dz_k) = \tilde{w}_k(z_{k-1})\Gamma_{k-1}(dz_{k-1})M_k(z_{k-1}, dz_k) .$$

Hence, (S1) is satisfied with $w_k(z_{k-1}, z_k) = \tilde{w}_k(z_{k-1})$. In particular, if for all $k \in \{0, \dots, K\}$, $\Gamma_k(z_k) = \gamma_k(z_k)dz_k$, where γ_k is a positive p.d.f., then $\tilde{w}_k(z_k) = \gamma_k(z_k)/\gamma_{k-1}(z_{k-1})$.

3. Assume that for $k \in \{1, \dots, K\}$, M_k is reversible w.r.t. Γ_k , i.e. $\Gamma_k(dz_{k-1})M_k(z_{k-1}, dz_k) = \Gamma_k(dz_k)M_k(z_k, dz_{k-1})$, and that there exists a positive measurable function such that $\Gamma_k(dz_{k-1}) = \tilde{w}_k(z_{k-1})\Gamma_{k-1}(dz_{k-1})$. Then, setting $L_{k-1} = M_k$, (S1) is satisfied.

C.2. Proof of (14)

For $k \in \{1, \dots, K\}$, $z_{k-1} \in \mathbb{R}^d$, denote by $G_{k, z_{k-1}}$ the mapping $u_k \mapsto T_{u_k}(z_{k-1})$. Our derivation below rely on the fact that for $k \in \{1, \dots, K\}$, $z_{k-1} \in \mathbb{R}^d$, $G_{k, z_{k-1}}$ is a C^1 -diffeomorphism. This is the case for the Langevin mappings. Note, similarly to the density considered in Section 3, that $m_k(z_{k-1}, z_k) = \varphi(G_{k, z_{k-1}}^{-1}(z_k))J_{G_{k, z_{k-1}}^{-1}}(z_k)$. When $K = 1$, we have

$$\begin{aligned} \int \log(w_1(z_0, z_1))q_\phi^1(z_{0:1} | x)dz_{0:1} &= \int \log(w_1(z_0, z_1))q_\phi(z_0 | x)J_{G_{1, z_0}^{-1}}(z_1)\varphi(G_{1, z_0}^{-1}(z_1))dz_{0:1} \\ &= \int \log(w_1(z_0, T_{1, u_1}(z_0)))q_\phi(z_0 | x)\varphi(u_1)dz_0du_1 , \end{aligned}$$

where we have performed the change of variables $u_1 = G_{1, z_0}^{-1}(z_1)$, hence $z_1 = G_{1, z_0}(u_1) = T_{1, u_1}(z_0)$. Let now K be in \mathbb{N}^* . In general, we write

$$\begin{aligned} \mathcal{L}_{\text{SIS}} &= \int \log \left(\prod_{k=1}^K w_k(z_{k-1}, z_k) \right) q_\phi^K(z_{0:K} | x) dz_{0:K} = \int \log \left(\prod_{k=1}^K w_k(z_{k-1}, z_k) \right) q_\phi(z_0 | x) \prod_{k=1}^K m_k(z_{k-1}, z_k) dz_{0:K-1} dz_K \\ &= \int \log \left(\prod_{k=1}^K w_k(z_{k-1}, z_k) \right) q_\phi(z_0 | x) \prod_{k=1}^{K-1} m_k(z_{k-1}, z_k) \varphi(G_{K, z_{K-1}}^{-1}(z_K)) J_{G_{K, z_{K-1}}^{-1}}(z_K) dz_{0:K-1} dz_K \\ &= \int \log \left(\prod_{k=1}^{K-1} w_k(z_{k-1}, z_k) w_K(z_{K-1}, T_{u_K}(z_{K-1})) q_\phi(z_0 | x) \right) \prod_{k=1}^{K-1} m_k(z_{k-1}, z_k) \varphi(u_K) dz_{0:K-1} du_K \end{aligned}$$

using the change of variables $u_K = G_{K, z_{K-1}}^{-1}(z_K)$. By an immediate backwards induction, we write

$$\mathcal{L}_{\text{SIS}} = \int \log \left(\prod_{k=1}^K w_k \left(\bigcirc_{i=1}^{k-1} T_{i, u_i}(z_0), \bigcirc_{i=1}^k T_{i, u_i}(z_0) \right) \right) q_\phi(z_0 | x) \varphi(u_{1:K}) dz_0 du_{1:K} .$$

C.3. Proof of Lemma 1

Let $\eta < L^{-1}$ and $u \in \mathbb{R}^D$. First we show that T_u^{MALA} is invertible. Consider, for each $(y, u) \in \mathbb{R}^{2d}$, the mapping $H_{y, u}(z) = y - \sqrt{2\eta}u - \eta \nabla \log \pi(z)$. We have, for $z_1, z_2 \in \mathbb{R}^d$,

$$\|H_{y, u}(z_1) - H_{y, u}(z_2)\| \leq \eta \|\nabla \log \pi(z_1) - \nabla \log \pi(z_2)\| \leq \eta L \|z_1 - z_2\|$$

and $\eta L < 1$. Hence $H_{y, u}$ is a contraction mapping and thus has a unique fixed point $z_{y, u}$. Hence, for all $(y, u) \in \mathbb{R}^{2d}$ there exists a unique $z_{y, u}$ satisfying

$$H_{y, u}(z_{y, u}) = z_{y, u} \Rightarrow y = z_{y, u} + \eta \nabla \log \pi(z_{y, u}) + \sqrt{2\eta}u = T_u^{\text{MALA}}(z_{y, u}).$$

This establishes the invertibility of T_u^{MALA} . The fact that the inverse of T_u^{MALA} is C^1 follows from a simple application of the local inverse function theorem.

D. ELBO AIS

D.1. Construction of the control variates

We prove in this section that the variance reduced objective we consider is valid. Sample now n samples $u_{0:K}^{1:n} \stackrel{\text{i.i.d.}}{\sim} \varphi_{d,K+1}$. For an index $i \in \{1, \dots, n\}$, given the initial point $z_0^i = V_{\phi,x}(u_0^i)$ and the innovation noise $u_{1:K}^i$, we sample the A/R booleans $a_{1:K}^i$. We introduce, in the main text, for $i \in \{1, \dots, n\}$

$$\tilde{W}_{n,i} = \frac{1}{n-1} \sum_{j \neq i} W(V_{\phi,x}(u_0^j), a_{1:K}^j, u_{1:K}^j),$$

\tilde{W}_i provides a reasonable estimate of the AIS ELBO but is independent from the i -th trajectory. We use this quantity as a control variate to reduce the variance of our gradient estimator by introducing

$$\begin{aligned} \widehat{\nabla \mathcal{L}_{\text{AIS}n}} &= n^{-1} \sum_{i=1}^n \nabla W(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) \\ &\quad + n^{-1} \sum_{i=1}^n \left[W(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) - \tilde{W}_{n,i} \right] \\ &\quad \times \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i). \end{aligned} \quad (\text{S4})$$

Proving its unbiasedness boils down to proving that the term $n^{-1} \sum_{i=1}^n \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i)$ has expectation zero. Let us compute for $i \in \{1 \dots, n\}$,

$$\begin{aligned} &\int \sum_{a_{1:K}^i} \varphi_{d,K+1}(u_{0:K}^i) A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) du_{0:K}^i = \\ &\quad \int \sum_{a_{1:K-1}^i} \varphi_{d,K+1}(u_{0:K}^i) \prod_{k=1}^K \alpha_{k,u_k^i}^{a_k^i}(z_{k-1}^i) \tilde{W}_{n,i} \left[\nabla \sum_{k=1}^{K-1} \log \alpha_{k,u_k^i}^{a_k^i}(z_{k-1}^i) + \sum_{a_K^i} \nabla \log \alpha_{K,u_K^i}^{a_K^i}(z_{K-1}^i) \right] du_{0:K}^i, \end{aligned}$$

denoting $z_0^j = V_{\phi,x}(u_0^j)$, $z_k^j = \mathcal{O}_{i,u_i^j}^{a_i^j}(z_0^j)$ by simplicity of notation. Yet, $\sum_{a_K^i} \alpha_{K,u_K^i}^{a_K^i}(z_{K-1}^i) = 1$ exactly, thus $\sum_{a_K^i} \alpha_{K,u_K^i}^{a_K^i}(z_{K-1}^i) \nabla \log \alpha_{K,u_K^i}^{a_K^i}(z_{K-1}^i) = 0$. We can thus show by an immediate induction that $\int \sum_{a_{1:K}^i} \varphi_{d,K+1}(u_{0:K}^i) \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) du_{0:K}^i = 0$, as $\tilde{W}_{n,i}$ is a constant in that integral by independence of the samples for $i \in \{1 \dots, n\}$. Moreover, as

$$\begin{aligned} &\int \sum_{a_{1:K}^{1:n}} \sum_{i=1}^n \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) \prod_{\ell=1}^n \varphi_{d,K+1}(u_{0:K}^\ell) du_{0:K}^{1:n} = \\ &\quad \int \sum_{i=1}^n \sum_{a_{1:K}^{-i}} \left[\int \sum_{a_{1:K}^i} \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i) \varphi_{d,K+1}(u_{0:K}^i) du_{0:K}^i \right] \prod_{\ell \neq i} \varphi_{d,K+1}(u_{0:K}^\ell) du_{1:K}^{-i}, \end{aligned}$$

then $n^{-1} \sum_{i=1}^n \tilde{W}_{n,i} \nabla \log A(V_{\phi,x}(u_0^i), a_{1:K}^i, u_{1:K}^i)$ is of zero expectation, and (S4) is an unbiased estimator of the gradient.

D.2. Discussion of (Wu et al., 2020)

In (Wu et al., 2020), authors consider a MCMC VAE inspired by AIS. The model used however is quite different in spirit to what is performed in this work. (Wu et al., 2020) use Langevin mappings and accept/reject steps in their VAE. Note however that the A/R probabilities defined are written as

$$\alpha(x, y) = 1 \wedge \pi(y)/\pi(x),$$

different from (21). Moreover, even though accept/reject steps are considered, the score function estimator (25) is not taken into account.

Finally, the initial density of the sequence is not taken to be some variational mean field initialization but directly the prior in the latent space. As a result, the scores obtained by the MCMC VAE are less competitive than that of the RNVP VAE presented in (Wu et al., 2020, Table 3.), contrary to what is presented here.