

Optimization of Annealed Importance Sampling Hyperparameters

Shirin Goshtasbpour^{1,2}[0000–0002–4902–165X]
and Fernando Perez-Cruz^{1,2}[0000–0001–8996–5076]

¹ Computer Science dep., ETH Zurich, Ramistrasse 101, 8092 Zurich, Switzerland

² Swiss Data Science Center, Turnerstrasse 1, 8092 Zurich, Switzerland
`shirin.goshtasbpour@inf.ethz.ch`

Abstract. Annealed Importance Sampling (AIS) is a popular algorithm used to estimate the intractable marginal likelihood of deep generative models. Although AIS is guaranteed to provide unbiased estimate for any set of hyperparameters, the common implementations rely on simple heuristics such as the geometric average bridging distributions between initial and the target distribution which affect the estimation performance when the computation budget is limited. Optimization of fully parametric AIS remains challenging due to the use of Metropolis-Hasting (MH) correction steps in Markov transitions. We present a parametric AIS process with flexible intermediary distributions and optimize the bridging distributions to use fewer number of steps for sampling. A reparameterization method that allows us to optimize the distribution sequence and the parameters of Markov transitions is used which is applicable to a large class of Markov Kernels with MH correction. We assess the performance of our optimized AIS for marginal likelihood estimation of deep generative models and compare it to other estimators.

Keywords: Annealed Importance Sampling · Partition function estimation · Generative models

1 Introduction

Deep generative models, developed over a decade, are now capable of simulating complex distributions in high dimensional space and synthesizing high quality samples in various domains such as natural images, text and medical data. Many of these models are built with the assumption that data $x \in \mathcal{X}$ resides close to a low dimension manifold. In this case, data can be represented using latent variable z drawn from prior distribution on \mathbb{R}^d with density $p(z)$ and the observation model $p(x|z) = p(x|f_\theta(z))$ is given by a parametric mapping $f_\theta : \mathbb{R}^d \rightarrow \Gamma$ where Γ is the space of parameters of the likelihood $p(x|z)$.

For complex f_θ mappings, evaluation of log marginal likelihood

$$\log p(x) = \log \int p(z)p(x|z)dz$$

is intractable and the posterior density $p(z|x) \propto p(x|z)p(z)$ is known only up to a constant normalization factor³. In this case, latent variable inference, probabilistic model evaluation and model comparison are performed using variational approximation of the posterior [3,24,38] or via sampling methods like Markov Chain Monte Carlo (MCMC) [4], Nested sampling [43,32], Sequential Monte Carlo (SMC) [8] and Annealed Importance Sampling (AIS) [34]. In Variational Inference (VI) the posterior is approximated with the most similar probability density from the family of distributions $\mathcal{Q} = \{q_\phi(z|x) : \phi \in \Phi\}$ indexed by parameter ϕ . Success of VI depends on sufficient expressivity of members of \mathcal{Q} and our ability to find the optimal member of this family by, for instance, maximizing the Evidence Lower Bound (ELBO)

$$\mathcal{L}(\phi) = \int \log \left(\frac{p(x, z)}{q_\phi(z|x)} \right) q_\phi(z|x) dz \leq \log p(x) \quad (1)$$

where the equality happens only when $q_\phi(z|x) = p(z|x) \in \mathcal{Q}$.

In MCMC we use Markov kernels with unique stationary distribution $p(z|x)$ and sample a Markov chain $(z_k)_{k \in [M]}$ by iterative application of this kernel on an initial particle $z_0 \sim q_\phi(z|x)$. Contrary to parametric VI, under mild assumptions, it is theoretically guaranteed that distribution of z_M converges to the target distribution $p(z|x)$ as M goes to infinity. However, if the posterior has multiple modes or heavy tails, convergence can require large number of iterations and therefore, be computationally prohibitive. Among the sampling methods, SMC and AIS are of particular interest as they produce unbiased estimation of marginal likelihood regardless of the computational budget by assigning importance weights to samples [40,50,31]. In these algorithms an auxiliary sequence of distributions, $(\pi_k)_{k \in [M]}$, is used to bridge a simple proposal distribution with density $\pi_0(z) = q_\phi(z|x)$ and the target density $\pi_M(z) = p(z|x)$. This sequence is defined via unnormalized densities $(\gamma_k)_{k \in [M]}$ where $\pi_k(z) = \gamma_k(z)/Z_k$ and $\gamma_0(z) = \pi_0(z)$ and $\gamma_M(z) = p(x, z)$. The algorithm produces N Markov chain samples $(z_k^j)_{k \in [M]}$ for $j \in [N]$ and their corresponding importance weights denoted by w^j as follows: Initially z_0^j is sampled from γ_0 . Then z_k^j is approximately sampled from γ_k using a Markov kernel (typically a gaussian distribution around z_{k-1}^j or a transition with invariant distribution γ_k)⁴. The marginal likelihood is approximated with Monte Carlo method

$$p(x) = \mathbb{E}[w] \approx \hat{p}^N(x) := \frac{1}{N} \sum_{j \in [N]} w^j \quad (2)$$

³ We use normalization factor and marginal likelihood $p(x)$ interchangeably in this paper

⁴ SMC has an additional resampling step to draw exact samples from γ_k to reduce the variance of importance weight although it sometimes results in insufficient sample diversity

where the expectation is taken over joint distribution of the Markov chains. Variance of $\hat{p}^N(x)$ depends on the selected density sequence and hyperparameters especially when computation resources are scarce. However, little work is available on how to optimize the intermediary distributions. In this paper we focus on optimization of AIS hyperparameters for more accurate estimation of log marginal likelihood with only a few annealing distributions. We optimize the sequence of distributions between proposal and target distributions as well as the hyperparameters of Markov kernels. We have made the following contributions:

- We propose the parameterization of a continuous time density path between the proposal and target distributions which can define general density functions. Contrary to [24,25] the densities used do not need to be normalized and we don't require their exact samples since we use sampling algorithms to gradually transition between the intermediary distributions.
- To optimize the bridging distributions we minimize Jefferys and inverse KL divergences between the distributions of AIS process and its auxiliary reverse process (defined in Section 2.1). Inspired from recent developments in Monte Carlo VI (MCVI) methods, we reparameterize particle states with a mixture of Normalizing Flows and derive a low variance consistent estimation of the gradient of our optimization objective.
- We show that Jefferys divergence captures both bias and variance of the estimation while optimization of inverse KL divergence reduces the bias of estimation. We further evaluate our method on deep generative models with different training procedures.

The combination of both procedures allows reducing the variance of the importance sampling weights providing more accurate estimates with fewer particles and intermediary distributions. The rest of this paper is organized as follows: Section 2 is dedicated to reviewing vanilla AIS algorithm. In Section 3 we present our parameterization of AIS process which results in flexible bridging distributions and motivate our optimization objective. In Section 4 we illustrate a reparameterization method and derivation of the objective gradient estimates. Finally, we analyze the accuracy of marginal likelihood estimation and its variance on synthetic and image data in Section 6.

2 Background

In this Section we give a brief introduction to AIS algorithm and its popular adaptive versions. For the rest of this section, we assume that the observation x is fixed and $\pi(z) = p(z|x)$ is the target density function which we can evaluate up to a normalizing constant $\tilde{\pi}(z) = p(z, x)$ where $\pi(z) = \tilde{\pi}(z)/Z$ and $Z = p(x)$. We also define a proposal distribution with normalized density $\gamma_0(z)$ which is easy to sample and evaluate, such as the variational posterior $\gamma_0(z) = q_\phi(z|x)$.

2.1 Recap of Annealed Importance Sampling

To sample from the target distribution, AIS algorithm requires a sequence of distributions which change from the proposal distribution to the target. To

match with our parameterization of the bridging distributions, we consider the generalization of this sequence to a continuously indexed path of density functions. With some abuse of notation, let $\gamma : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}$ with mapping $(t, z) \mapsto \gamma(t, z)$ denote a path of density functions indexed by $t \in [0, 1]$, with fixed endpoints $\gamma(0, \cdot) = q_\phi(\cdot|x)$ and $\gamma(1, \cdot) = \tilde{\pi}(\cdot)$. Assume the density functions have non-increasing support (i.e. $\text{supp}(\gamma(t', \cdot)) \subseteq \text{supp}(\gamma(t, \cdot))$ for $0 \leq t \leq t' \leq 1$). Let $t_k = k/M$ for $k \in \{0, \dots, M\}$ be the linear schedule in $[0, 1]$ with which the path of density functions is discretized and let $\gamma_k(\cdot) = \gamma(t_k, \cdot)$. We assume that the $\gamma(\cdot, z)$ curves are flexible enough to absorb any increasing mapping $\beta : [0, 1] \rightarrow [0, 1]$ (representing a discretization of $[0, 1]$ at $(\beta(t_k))_{k \in \{0, \dots, M\}}$).

AIS algorithm starts by sampling the initial particles $z_0 \sim \gamma_0(z)$ and evaluating the initial weights $w_0 = 1$. Then, particles are gradually moved in iteration k using Markov kernel $\vec{T}_k(z_k|z_{k-1})$ such that the distribution of particles z_M approaches the target distribution $\pi(z)$. The particle trajectory distribution in AIS has the joint density $\vec{q}_{\phi, M}(z_{0:M}) = \gamma_0(z_0) \prod_k \vec{T}_k(z_k|z_{k-1})$. In each iteration a second Markov kernel $\overleftarrow{T}_k(z_{k-1}|z_k)$ is used to approximate the probability density of backward transition $z_k \rightarrow z_{k-1}$ such that the ratio $w_k(z_{k-1}, z_k) = \gamma_k(z_k) \overleftarrow{T}_k(z_{k-1}|z_k) / \gamma_{k-1}(z_{k-1}) \vec{T}_k(z_k|z_{k-1})$ is well-defined on \mathbb{R}^d and the importance weights are adjusted by multiplying $w_k(z_{k-1}, z_k)$ with the current weights. The final importance weights can be evaluated as

$$w(z_{0:M}) = \frac{\tilde{\pi}(z_M) \overleftarrow{q}_{\phi, M}(z_{0:M-1}|z_M)}{\vec{q}_{\phi, M}(z_{0:M})} = \prod_{k=1}^M w_k(z_{k-1}, z_k)$$

where $\overleftarrow{q}_{\phi, M}(z_{0:M-1}|z_M) = \prod_k \overleftarrow{T}_k(z_{k-1}|z_k)$.

The optimal backward process $\overleftarrow{q}_{\phi, M}(z_{0:M-1}|z_M)$ with zero variance estimator $\hat{p}^N(x)$ is given by inverse kernels

$$\overleftarrow{T}_k^*(z_{k-1}|z_k) = \frac{\vec{q}_{\phi, k-1}(z_{k-1}) \vec{T}_k(z_k|z_{k-1})}{\vec{q}_{\phi, k}(z_k)}$$

where $\vec{q}_{\phi, k}(z_k)$ is the marginalization of $\vec{q}_{\phi, M}(z_{0:M})$ over all z_m except $m \neq k$. However it is intractable to use $\overleftarrow{T}_k^*(z_{k-1}|z_k)$ for sampling [8].

Alternatively, we can use a heuristically fixed $(\gamma_k)_k$ to guide the samples towards the target's basins of energy. One may choose \vec{T}_k from popular MCMC kernels like Random Walk Metropolis-Hastings (RWMH), Hamiltonian Monte Carlo (HMC), Metropolis Adjusted Langevin Algorithm (MALA) or Unadjusted Langevin Algorithm (ULA) with π_k as their stationary distribution [45,46]. MH-corrected kernels typically do not admit a transition probability density. However, due to detailed balance $\gamma_k(z_k) \vec{T}_k(z_{k-1}|z_k) = \gamma_k(z_{k-1}) \vec{T}_k(z_k|z_{k-1})$ the kernel \vec{T}_k is reversible with respect to γ_k and with backward kernel set to its reversal $\gamma_k(z_{k-1}) \vec{T}_k(z_k|z_{k-1}) / \gamma_k(z_k)$ we can obtain well-defined importance weight updates $w_k(z_{k-1}, z_k) = \gamma_k(z_{k-1}) / \gamma_{k-1}(z_{k-1})$ [34]. When \vec{T}_k is not γ_k -invariant, an approximate reversal kernel with same assumption may be constructed from

$\overleftarrow{\mathcal{T}}_k(z_{k-1}|z_k) = \overrightarrow{\mathcal{T}}_k(z_{k-1}|z_k)$ while the weight updates preserves their original form [46].

In another approach, reparameterizable forward and backward transition kernels can be used where the parameters are optimized to push samples from one predetermined bridging distribution to the next, which limits the complexity of the transition probability density considerably [31,49,1].

3 Optimized Path Annealed Importance Sampling

Typical implementations of AIS use predefined heuristics for the sequence $(\gamma_k)_k$ and the schedule β . Popular choices include the geometric average density path $\gamma(t, z) = \tilde{\pi}(z)^{\beta(t)} q_\phi(z|x)^{\beta(1-t)}$ with linear $\beta(t) = t$ or geometric $\beta(t) = \alpha^{\log t}$ schedule, for some real $\alpha > 1$. In contrast, we propose to optimize the bridging distribution path $\gamma(t, \cdot)$ in a class of positive parametric functions. If the class is sufficiently large the schedule β will be implicit in γ and therefore we can fix the schedule to be linear and only focus on parameterizing the $\gamma(t, \cdot)$. We use Deep Neural Networks (DNN) to define our parametric density function path. We name our procedure Optimized Path Annealed Importance Sampling (OP-AIS).

3.1 Parameterized AIS

We opt to parameterize the time dependent path γ as opposed to parameterizing the intermediary distributions of each iteration individually. The main reason for this choice is that it results in a density path that is continuous in t and allows parameters to be shared by intermediate distributions. We use a continuous mapping $u_\phi : \mathcal{X} \times \mathcal{Z} \times [0, 1] \rightarrow \mathbb{R}$ with additional terms that adjust the boundary constraints at $t \in \{0, 1\}$. We consider the following boundary adjustment which coincides with geometric average of the proposal and target distributions when $u_\phi \equiv 0$.

$$\begin{aligned} \log \gamma(t, z) = & u_\phi(x, z, t) \\ & + (1 - t) [-u_\phi(x, z, 0) + \log \gamma_0(z)] \\ & + t [-u_\phi(x, z, 1) + \log \tilde{\pi}(z)] \end{aligned} \quad (3)$$

If $u_\phi(x, z, t) \in (-\infty, \infty)$ is well-defined for all $x \in \mathcal{X}$, $z \in \mathbb{R}^d$ and $t \in (0, 1)$ and the proposal is chosen such that $\text{supp}(\gamma(1, \cdot)) \subseteq \text{supp}(\gamma(0, \cdot))$ then $\text{supp}(\gamma(1, \cdot)) = \text{supp}(\gamma(t, \cdot))$ for all $t \in [0, 1]$ and the importance weights updates are always defined using the common gaussian or invariant transition kernels. Notably, we can apply this parameterization to any arbitrary path of density functions by twisting the path with multiplication of an arbitrary positive function $\exp(-u_\phi(x, z, t))$ with correct boundary behavior.

3.2 Optimization for the parametric distribution

Given an unnormalized target distribution, Effective Sample Size (ESS) and its conditional extension are the de facto measure used to adaptively tune AIS

[8,21,22,34]. ESS of N parallel runs of AIS is given by

$$\text{ESS}(w^1, \dots, w^N) = \frac{(\sum w^j)^2}{N \sum (w^j)^2}$$

and is expected to reflect the estimator variance odds with exact samples from $\pi(z)$ and importance sampling (2). ESS is sensitive to weight degeneracy problem that is prominent in AIS and other sequential importance sampling methods (when only a few particles contribute to the Monte Carlo approximate). However, high ESS does not determine if the particles have sufficient dispersion according to the target distribution. See [11] for further discussion on ESS and its defects.

Grosse et. al. in [16] derived a bound on $\log Z$ by running the AIS algorithm in forward and backward directions. Original direction of AIS evaluates $\log \vec{w} = \log w(z_{0:M})$ which in expectation lower bounds $\log Z$.

$$\mathbb{E}_{\vec{q}_{\phi,M}}[\log \vec{w}] \leq \log Z \quad (4)$$

Drawing sample trajectories $z'_{0:M}$ from the reverse direction of AIS algorithm defined by the ancestral sampling in joint distribution $\pi(z'_M) \overleftarrow{q}_{\phi,M}(z'_{0:M-1}|z'_M)$ we can compute random variable $\log \overleftarrow{w} = \log w(z'_{0:M})$ which upper bounds $\log Z$ in expectation. The gap between the two expectations in forward and backward processes is called the BiDirectional Monte Carlo (BDMC) gap and is equivalent to Jefferys divergence between these distributions.

$$\begin{aligned} \mathcal{L}^{\text{BDMC}}(x) &= \mathbb{E}_{\pi \overleftarrow{q}_{\phi,M}}[\log \overleftarrow{w}] - \mathbb{E}_{\vec{q}_{\phi,M}}[\log \vec{w}] \\ &= \text{D}_{\text{KL}}(\pi \overleftarrow{q}_{\phi,M} || \vec{q}_{\phi,M}) + \text{D}_{\text{KL}}(\vec{q}_{\phi,M} || \pi \overleftarrow{q}_{\phi,M}) \end{aligned} \quad (5)$$

This bound is also frequently used to assess the accuracy of the AIS estimator [15,50,20]. Unbiased estimation of this bound requires exact samples from the target distribution. Therefore, an approximation is made replacing the data samples with the synthesized samples by the generative model, where we have access to an underlying latent variable. This approximation results in more accurate log-likelihood estimation of samples closer to the model which affects the objectivity of the test. Assuming a long chain and close to perfect transitions the final particles from AIS forward process can be used to approximate this bound instead. However, with low computation budget this marginal distribution may be far off from the target especially in the initial optimization iterations. Although, the BDMC gap achieves a natural trade-off between minimization of the bias in ELBO bound and log empirical variance of the importance weights (Figure 2), the two terms in $\mathcal{L}^{\text{BDMC}}(x)$ have contradicting gradients resulting in unstable optimization of hyperparameters with gradient based methods.

In order to achieve low variance importance weights we need to match the distributions of the forward and backward processes. It is common to minimize the inverse KL-divergence between these distributions similar to maximization of (1) as this results in reducing the bias in (4).

$$\mathcal{L}^{\text{AIS}}(x) = \text{D}_{\text{KL}}(\vec{q}_{\phi,M} || \tilde{\pi} \overleftarrow{q}_{\phi,M}) = -\mathbb{E}_{\vec{q}_{\phi,M}}[\log \vec{w}]$$

We can evaluate an unbiased estimate of inverse KL divergence using samples from $\vec{q}_{\phi, M}$.

As was shown by authors of [9, 28] under some assumptions, the variance of importance weights can be controlled by the inverse KL-divergence between two distributions. In particular, with some adjustment to the AIS setting we have

$$D_{\text{KL}}(\vec{q}_{\phi, M} || \tilde{\pi} \overleftarrow{q}_{\phi, M}) \approx \frac{1}{2Z^2} \text{Var}_{\vec{q}_{\phi, M}}[w(z_{0:M})]$$

Therefore, by minimizing the inverse KL-divergence we minimize the log marginal likelihood bias and the variance of the importance weights, simultaneously. The equality holds when $\pi(z_k) \overleftarrow{q}_{\phi, M}(z_{0:M-1} | z_M) = \vec{q}_{\phi, M}(z_{0:M})$ which results in zero importance weight variance and $\mathcal{L}^{\text{AIS}}(x) = 0$.

4 Stochastic Optimization for OP-AIS

In order to evaluate the gradients of $\mathcal{L}^{\text{AIS}}(x)$ with respect to the parameters ϕ we use a combination of reparameterization trick and score function estimation introduced by [42] and [31]. The MH-corrected Markov kernel $\vec{T}_k(z_k | z_{k-1})$ is given by

$$\begin{aligned} \vec{T}_k(z_k | z_{k-1}) &= r(z_k | z_{k-1}) \alpha_k(z_{k-1}, z_k) \\ &+ \int (1 - \alpha_k(z | z_{k-1})) r(z | z_{k-1}) dz \delta_{z_{k-1}}(z_k) \end{aligned}$$

with $r(\cdot | z_{k-1})$ denoting the conditional probability density of proposed particle in the k th iteration, $\alpha_k(z_{k-1}, z) = \left(1 \wedge \frac{\gamma_k(z) r(z_{k-1} | z)}{\gamma_k(z_{k-1}) r(z | z_{k-1})}\right)$ the acceptance rate at this iteration and δ_z denoting the Dirac measure at z . We assume that the proposal $r(z | z_{k-1})$ is reparameterizable with a auxiliary random variable $\epsilon_k \sim \eta(\epsilon)$ and transformation $z = T_k(\epsilon_k, z_{k-1})$ which is surjective in second argument. After transformation the proposal state is accepted with probability $\alpha_k(z_{k-1}, T_k(\epsilon_k, z_{k-1}))$ (denoted with α_k as a shorthand) and we set $z_k = T_k(\epsilon_k, z_{k-1})$. Otherwise, we keep the old state and set $z_k = z_{k-1}$. We determine the proper action by sampling a binary random variable $a_k \sim \text{Bern}(\alpha_k(z_{k-1}, T_k(\epsilon_k, z_{k-1})))$.

For instance, for a Random Walk kernel \vec{T}_k , we assume that $\eta(\epsilon)$ is the probability density function of standard normal distribution and $T_k(\epsilon, z) = z + \Sigma^{1/2} \epsilon$ for a positive definite covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$. We refer the readers to the derivation of transformations corresponding to HMC, MALA and ULA by [45] for further details.

Using the above reparameterization we get

$$\mathcal{L}^{\text{AIS}}(x) = -\mathbb{E}_{\gamma_0(z_0) \prod_k \eta(\epsilon_k) \prod_k \zeta_k(a_k | z_0, \epsilon_{1:k}, a_{1:k-1})} [\log w]$$

$$\text{for } \zeta_k(a^k | z_0, \epsilon_{1:k}, a_{1:k-1}) = \alpha_k^{a_k} (1 - \alpha_k)^{1-a_k}.$$

The gradient of $\mathcal{L}^{\text{AIS}}(x)$ is given by

$$\begin{aligned}\nabla_{\phi}\mathcal{L}^{\text{AIS}}(x) &= -\mathbb{E}[\nabla_{\phi}\log w] \\ &\quad - \sum_k \mathbb{E}[\log w \nabla_{\phi}\log \zeta_k(a_k|z_0, \epsilon_{1:k}, a_{1:k-1})]\end{aligned}$$

We can derive $\mathcal{L}^{\text{AIS}}(x)$ approximation using the particle and importance weights generated from the AIS .

$$\begin{aligned}\nabla_{\phi}\hat{\mathcal{L}}^{\text{AIS}}(x) &= -\frac{1}{N}\sum_{i=1}^N \nabla_{\phi}\log w^i \\ &\quad + \underbrace{-\frac{1}{N}\sum_{i=1}^N \sum_k \log w^i \nabla_{\phi}\log \zeta(a_k^i|z_0^i, \epsilon_{1:k}^i, a_{1:k-1}^i)}_{\nabla_{\phi}\hat{\mathcal{L}}^{\text{score}}(x)}\end{aligned}$$

$\nabla_{\phi}\hat{\mathcal{L}}^{\text{AIS}}(x)$ is a strongly consistent estimator of $\nabla_{\phi}\mathcal{L}^{\text{AIS}}(x)$ [27].

Score function estimator is notorious for its high variance. It is a standard practice to omit $\nabla_{\phi}\mathcal{L}^{\text{score}}$ in parameter updates for more stable optimization. However, this results in biased gradient estimate. To reduce the high variance of $\nabla_{\phi}\hat{\mathcal{L}}^{\text{score}}(x)$ we can employ the variation reduction techniques [33,36]. For instance, in our experiments we use the common particles and weights to estimate the two expectations in $\mathcal{L}^{\text{AIS}}(x)$ and the one-out control variates proposed by [30] to substitute $\nabla_{\phi}\hat{\mathcal{L}}^{\text{score}}(x)$ with

$$\sum_{i=1}^N \left[\frac{NR^i - \sum_j R^j}{N-1} \right] \sum_k \nabla_{\phi}\log \zeta_k(a_k^i|z_0^i, \epsilon_{1:k}^i, a_{1:k-1}^i)$$

where $R^i = -\log w^i/N$.

Pseudocode of the optimization algorithm is given in Algorithm (1).

5 Literature Review

We can interpret AIS as the discretization of the path sampling algorithm where the thermodynamics integral identity holds

$$\log Z = \int_0^1 \mathbb{E}_{\gamma(t,z)/Z_t} \left[\frac{d}{dt} \log \gamma(t, z) \right] dt$$

for $Z_t = \int \gamma(t, z) dz$. Gelman et. al. in [13] derived the optimal path and scaling of the bridging distributions for minimum variance estimation. Evaluation of optimal path is intractable in general. In turn, authors of [6,2] provided closed form optimal path for special cases. Viewing the importance weight updates as the finite difference approximations, the authors in [17] derive the asymptotic

Algorithm 1: Optimization of parameterized AIS

```

input : Target  $\tilde{\pi}$ , proposal density  $\gamma_0$ , and  $\eta$ 
output: Parameters  $\phi$ 
1 for  $K$  epochs do
2   Set  $\Delta \leftarrow 0$ 
3   for  $i = 1$  to  $N$  do
4     Set  $\log w^i \leftarrow 0$ ,  $\log \zeta^i \leftarrow 0$ 
5     Draw  $z^i \sim q_0(z)$ 
6     Draw  $\epsilon_{1:M}^i \sim \prod_k \eta(\epsilon_k)$ 
7   for  $k = 1$  to  $M$  do
8     for  $i = 1$  to  $N$  do
9       Set  $\log w^i \leftarrow \log w^i + \log \frac{\gamma_k(z^i)}{\gamma_{k-1}(z^i)}$ 
10      Draw  $a_k^{(n)} \sim \text{Bern}(\alpha_k)$ 
11      Set  $\log \zeta^i \leftarrow \log \zeta^i + \log \zeta(a_k^i | z_0^i, \epsilon_{1:k}^i, a_{1:k-1}^i)$ 
12      Set  $z^i \leftarrow T_k(\epsilon_k^i, z^i)^{a_k^i} (z^i)^{1-a_k^i}$ 
13   Get  $R^i$  for  $i \in [N]$ 
14    $\phi \leftarrow \phi - \kappa \nabla_{\phi} \hat{\mathcal{L}}^{\text{AIS}}(x)$ 
    
```

bias of $\log w$ estimator and propose moment average path between members of exponential family. A variational representation of the annealing distributions minimizing weighted α -divergences between proposal and target distributions are derived in [17,29]. In [26], the authors minimize the asymptotic variance of estimator $\log w$ for a given parametric path of distributions.

The analysis in these works are derived based on the assumption of reaching equilibrium in each iteration and independence of particles. This is not the case when transition kernels are not invertible. On the other hand, exhaustive heuristic optimization to tune AIS hyperparameters is laborious [40]. Other methods to adapt AIS hyperparameters include monitoring the acceptance rate in Markov transitions to adjust its parameters [21,46] and designing problem specific or heuristic intermediary distribution sequences [21,40,29,17]. On the contrary, while our method requires training before evaluation, we can amortize the overhead by parameter sharing and reduce the computation cost by using only a small subset of validation samples. As a consequence, we find an efficient and shorter annealing process for sampling. Our aim is to enable AIS to achieve high accuracy despite limited computation resources.

Recently, Variational Inference and Monte Carlo methods have been combined to increase the flexibility of parametric variational family \mathcal{Q} and benefit from the convergence properties of sampling methods. Authors of [42,48,7,19,45] combine MCMC and VI by extending the space of latent variables $(z_k)_{k \in [M]}$. A framework with parametric transitions for SMC algorithms and a surrogate lower bound ELBO are provided in [18,47,31]. For example, in [28], the authors derive an

approximation of the asymptotic variance of importance weights in Variational Filtering Objectives, and in [49,1], flows are used to overcome energy density barriers. A differentiable alternative to MH-correcting step in HMC transitions are proposed [52], while in [12] the authors derive the importance weights resulting from unadjusted HMC process and its reversal.

Our work in particular bares more similarity to MCVAE proposed by [46] among others. In MCVAE the authors use AIS with heuristic density path for low variance estimation of the log marginal likelihood during training of a latent variable model. Our goal here is to evaluate an already trained model by optimizing the intermediary distributions and we argue that the learnt AIS parameters by MCVAE are not efficient since the model is constantly changing during training and the posterior approximator is known to overfit the model [50].

6 Experiments

Here, we evaluate the proposed algorithm on complex synthetic 2d distributions which are often used to benchmark generative models [37]. Target distributions are illustrated in Figure 1. We compare our algorithm to vanilla AIS [34] with geometric average path and geometric and linear schedule. We use RWMH, HMC and MALA transitions and normal proposal with learnable mean and diagonal covariance matrix. $M \in \{1, 2, 4, 8, 16, 64\}$ bridging distributions and 256 samples are used for training. u_ϕ is implemented with a DNN with one hidden layer with 4 dimensions and LeakyReLU(0.01) nonlinearity except at the output layer.

The parametric AIS is trained with Adam optimizer [23] with learning rate of 0.03 and betas = (0.9, 0.999) for 100 epochs. The RWMH kernels have local normal proposal steps with learnable diagonal covariance matrix and the step size in all transition kernels are trained along with other parameters. We use cross validation to choose the step size in vanilla AIS from $\{0.01, 0.05, 0.1, 0.5, 1.\}$. The normalization constants are estimated using $N = 4096$ samples. To act as the reference normalization constant we use vanilla AIS with $M = 1024$, 10 MCMC steps per iteration with 3 different step sizes. All the code is written with Pytorch [35] and experiments are run on a GeForce GTX 1080 Ti. Code is available here: https://github.com/shgoshtasb/op_ais.

In Figure 1 we illustrate particles generated from the samplers trained on each distribution with inverse KL divergence (PKL) and Jefferys divergence (PJ) color coded by their weight for $M = 8$. Warmer colors show particles with higher weights which are main contributors to the partition function estimation, while dark blue particles are less effective in the estimation. We observe that training with KL divergence results in smaller bias, while using Jefferys divergence objective results in more effective particles and lower variance estimation in Target 3-6. Figure 2 shows the log partition function estimation and the empirical variance of importance weights for the mentioned samplers. We index the target distributions in the same order they appear in the Figure 1. In most of the experiments, for small number of bridging distributions ($M \in \{2, 4, 8\}$) our method is able to

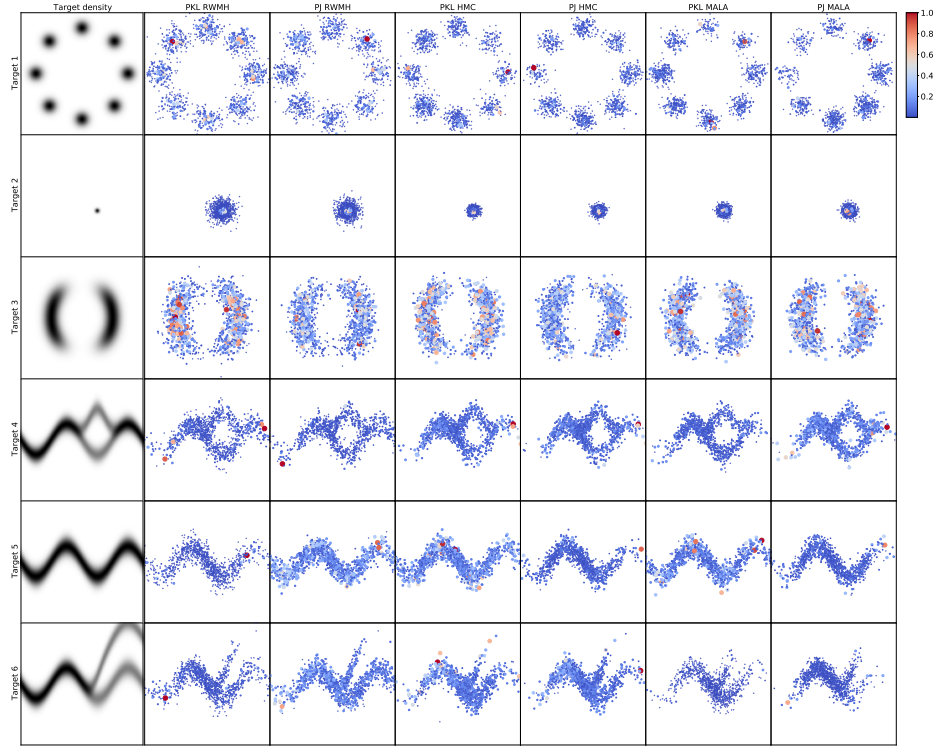


Fig. 1. Target distributions and particles of samplers with optimized path (trained with KL and Jefferys divergence) with RWMH, HMC and MALA transitions and $M = 8$. Particles with larger weights are coded to show comparable scale to maximum weight in each plot. Larger dots with warmer colors show particles with larger weights, with red for maximum weight in each set and dark blue for particles with 0.1 of the maximum weight

improve over vanilla AIS algorithm with linear schedule and achieves slightly tighter lower bound in comparison to geometric schedule version. In comparison to training with inverse KL divergence, Jefferys divergence improves the variance of importance weights. we speculate this to be related to higher correlation of Jefferys divergence with weight variances.

Interestingly, in multimodal and heavy tail distributions, empirical variance for small M remains below the variance of exhaustive AIS (VI $M=1024$ in dashed Turquoise) and grows with M . With only one MCMC step used in the experiments, particles ultimately cover a small area and don't reach the high density regions in the target distribution, resulting in low empirical weight variance despite the fact that the actual variance is much larger than our observation. This is specially problematic with the geometric schedule as the few bridging distributions are placed very close to the proposal and the particles are not encouraged to move far. In the low M setting as it is unlikely to observe high

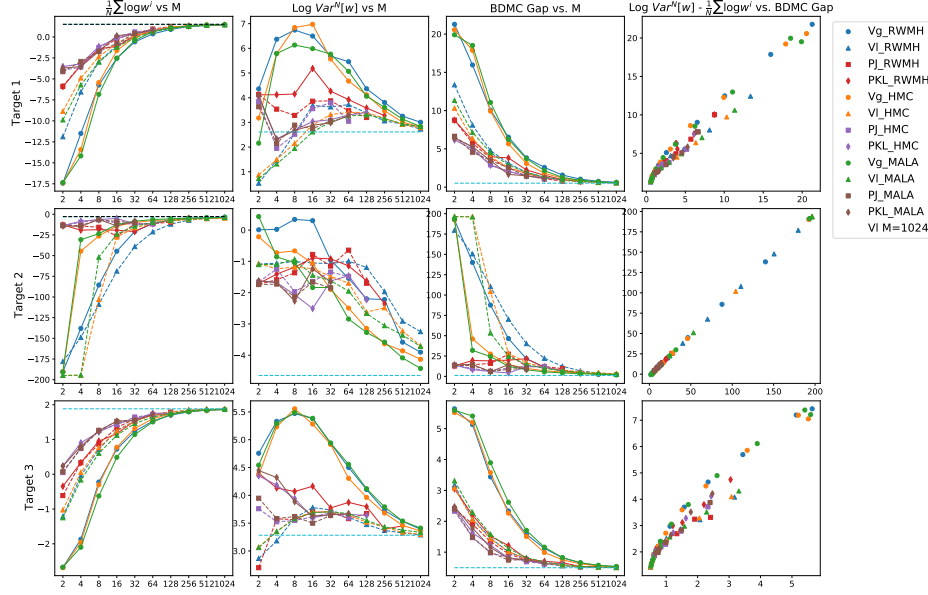



Fig. 2. From left to right: $\log Z$ estimation with Eq. 4, log empirical variance of importance weights and the Jefferys divergence of forward and backward processes vs M . (Vg, Vl for vanilla AIS with geometric and linear schedules, PJ and PKL for OP-AIS with Jefferys and KL-divergence optimization)

weight particles, it is misleading to use ESS or even empirical χ^2 -divergence to tune the hyperparameters, whereas the inverse KL divergence provides better guidance for tuning. Increasing M results in more dispersed particles and higher variance. For large enough M as the particles can reach high density regions in the target, we observe that this variance decreases with M . Variance in parametric samplers is more steady and has a small growth in comparison due to the optimized choice of bridging distributions for each M .

6.1 MNIST Dataset

In this section, we assess the effectiveness of our method on a DNN generative model trained on MNIST dataset. We use the binarized MNIST dataset with Bernoulli observation model with data dequantization and rescaling to $[0,1]$ [40]. The decoder model we consider is "Wide" architecture used in [50] with 50-1024-1024-1024-784 Fully-Connected (FC) layers with Tanh activations and Sigmoid output. This model is trained as a Variational AutoEncoder (VAE) [24], Importance Weighted AutoEncoder (IWAE) [5], Generative Adversarial Network (GAN) [14] and Bidirectional GAN (BiGAN) [10]. We compare OP-AIS with $M = \{2, 4, 8, 16\}$ and $N = 16$ particles per data samples HMC transition kernels with 1 leapfrog step to MCMC and Vanilla AIS with $M = 1024$ and HMC transitions with 10 leapfrog steps. OP-AIS is trained with Adam optimizer

**Fig. 3.** VAE sample reconstruction with posterior samples

with 0.001 learning rate and $\text{betas}=(0.5, 0.999)$ for 100 epochs on 1400 samples and validated on 600 samples. Reported log-likelihood values are over remaining 8000 test samples. Here, we use a symmetric encoder with Tanh activation and Dropout(0.2) at each hidden layer to transform x then input the output of that concatenated with the z to u_ϕ with 32-hidden units and LeakyReLU(0.01) activation. Table 1 shows the log marginal likelihood estimation using posterior

Table 1. Log likelihood estimation of generative models on binarized MNIST

Model	VI HMC	HMC	Encoder	PKL HMC	
				$M = 8$	$M = 16$
VAE	-87.8 ± 5.8	-521.3 ± 53.7	-95.7	-301.3 ± 36.0	-272.9 ± 35.0
IWAE	-85.8 ± 5.6	-563.4 ± 57.7	-88.5	-310.8 ± 35.5	-288.9 ± 35.4
GAN	-415.3 ± 48.8	-992.8 ± 87.4		-917.6 ± 77.5	-943.5 ± 80.0
BiGAN	-373.2 ± 44.6	-994.3 ± 85.5		-895.3 ± 81.6	-928.3 ± 83.9

samples from each sampling algorithm. While the training time is not sufficient to obtain values close to the encoder of the models in VAE or the exhaustive AIS, estimates are more accurate than HMC counterpart with similar computation time including training. We also observe that sample reconstruction with OP-AIS achieves similar samples to the test at least in style in Figure 3.

7 Conclusion

AIS yields an unbiased estimation for any annealing path, schedule and transition operator. However, the variance of the Monte Carlo estimator changes considerably with different choices. We propose an optimization procedure to

find the best set of intermediary distributions that matches the distribution of the AIS particle trajectory to joint distribution of backward process with target distribution as its marginal. We formulate AIS algorithm with MH-corrections as a mixture of NFs and minimize the the divergence between the two distributions.

It is important to mention that log-likelihood has to be used with caution for model comparison, since a model with very good or bad generative sample quality or memorization may have high log likelihood [44]. It is recommended to use other sample quality measures (e.g. [51,41]) and model comparison methods (e.g. [20,39]) along with AIS.

Acknowledgements This work is supported by funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 813999 for this project.

References

1. Arbel, M., Matthews, A.G., Doucet, A.: Annealed flow transport Monte Carlo. arXiv preprint arXiv:2102.07501 (2021)
2. Behrens, G., Friel, N., Hurn, M.: Tuning tempered transitions. *Statistics and computing* 22(1), 65–78 (2012)
3. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: A review for statisticians. *Journal of the American statistical Association* 112(518), 859–877 (2017)
4. Brooks, S., Gelman, A., Jones, G., Meng, X.L.: *Handbook of Markov chain Monte Carlo*. CRC press (2011)
5. Burda, Y., Grosse, R., Salakhutdinov, R.: Importance weighted autoencoders. arXiv preprint arXiv:1509.00519 (2015)
6. Calderhead, B., Girolami, M.: Estimating bayes factors via thermodynamic integration and population mcmc. *Computational Statistics & Data Analysis* 53(12), 4028–4045 (2009)
7. Caterini, A.L., Doucet, A., Sejdinovic, D.: Hamiltonian variational auto-encoder. arXiv preprint arXiv:1805.11328 (2018)
8. Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(3), 411–436 (2006)
9. Domke, J., Sheldon, D.: Importance weighting and variational inference. arXiv preprint arXiv:1808.09034 (2018)
10. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial feature learning. arXiv preprint arXiv:1605.09782 (2016)
11. Elvira, V., Martino, L., Robert, C.P.: Rethinking the effective sample size. arXiv preprint arXiv:1809.04129 (2018)
12. Geffner, T., Domke, J.: MCMC variational inference via uncorrected Hamiltonian annealing. arXiv preprint arXiv:2107.04150 (2021)
13. Gelman, A., Meng, X.L.: Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science* pp. 163–185 (1998)

14. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* 27 (2014)
15. Grosse, R.B., Ancha, S., Roy, D.M.: Measuring the reliability of mcmc inference with bidirectional Monte Carlo. *arXiv preprint arXiv:1606.02275* (2016)
16. Grosse, R.B., Ghahramani, Z., Adams, R.P.: Sandwiching the marginal likelihood using bidirectional Monte Carlo. *CoRR abs/1511.02543* (2015), <http://arxiv.org/abs/1511.02543>
17. Grosse, R.B., Maddison, C.J., Salakhutdinov, R.: Annealing between distributions by averaging moments. In: *NIPS*. pp. 2769–2777. Citeseer (2013)
18. Gu, S., Ghahramani, Z., Turner, R.E.: Neural adaptive sequential Monte Carlo. *arXiv preprint arXiv:1506.03338* (2015)
19. Hoffman, M.D.: Learning deep latent gaussian models with Markov chain Monte Carlo. In: *International Conference on Machine Learning*. pp. 1510–1519. PMLR (2017)
20. Huang, S., Makhzani, A., Cao, Y., Grosse, R.: Evaluating lossy compression rates of deep generative models. In: *International Conference on Machine Learning*. pp. 4444–4454. PMLR (2020)
21. Jasra, A., Stephens, D.A., Doucet, A., Tsagaris, T.: Inference for lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian Journal of Statistics* 38(1), 1–22 (2011)
22. Johansen, A.M., Aston, J.A., Zhou, Y.: Towards automatic model comparison: An adaptive sequential Monte Carlo approach (2015)
23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015), <http://arxiv.org/abs/1412.6980>
24. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
25. Kingma, D.P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., Welling, M.: Improved variational inference with inverse autoregressive flow. In: *Advances in Neural Information Processing Systems*. vol. 29 (2016)
26. Kiwaki, T.: Variational optimization of annealing schedules. *arXiv preprint arXiv:1502.05313* (2015)
27. Laubenfels, R.: Feynman–kac formulae: genealogical and interacting particle systems with applications (2005)
28. Maddison, C.J., Lawson, D., Tucker, G., Heess, N., Norouzi, M., Mnih, A., Doucet, A., Teh, Y.W.: Filtering variational objectives. *arXiv preprint arXiv:1705.09279* (2017)
29. Masrani, V., Brekelmans, R., Bui, T., Nielsen, F., Galstyan, A., Steeg, G.V., Wood, F.: q-paths: Generalizing the geometric annealing path using power means. *arXiv preprint arXiv:2107.00745* (2021)
30. Mnih, A., Rezende, D.: Variational inference for Monte Carlo objectives. In: *International Conference on Machine Learning*. pp. 2188–2196. PMLR (2016)
31. Naesseth, C., Linderman, S., Ranganath, R., Blei, D.: Variational sequential Monte Carlo. In: *International Conference on Artificial Intelligence and Statistics*. pp. 968–977. PMLR (2018)
32. Naesseth, C., Lindsten, F., Schon, T.: Nested sequential Monte Carlo methods. In: *International Conference on Machine Learning*. pp. 1292–1301. PMLR (2015)

33. Naesseth, C., Ruiz, F., Linderman, S., Blei, D.: Reparameterization gradients through acceptance-rejection sampling algorithms. In: Artificial Intelligence and Statistics. pp. 489–498. PMLR (2017)
34. Neal, R.M.: Annealed importance sampling. *Statistics and computing* 11(2), 125–139 (2001)
35. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
36. Ranganath, R., Gerrish, S., Blei, D.: Black box variational inference. In: Artificial intelligence and statistics. pp. 814–822. PMLR (2014)
37. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: International Conference on Machine Learning. pp. 1530–1538. PMLR (2015)
38. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: International conference on machine learning. pp. 1278–1286. PMLR (2014)
39. Sajjadi, M.S., Bachem, O., Lucic, M., Bousquet, O., Gelly, S.: Assessing generative models via precision and recall. *Advances in Neural Information Processing Systems* 31 (2018)
40. Salakhutdinov, R., Murray, I.: On the quantitative analysis of deep belief networks. In: Proceedings of the 25th international conference on Machine learning. pp. 872–879 (2008)
41. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* 29 (2016)
42. Salimans, T., Kingma, D., Welling, M.: Markov chain Monte Carlo and variational inference: Bridging the gap. In: International Conference on Machine Learning. pp. 1218–1226. PMLR (2015)
43. Skilling, J., et al.: Nested sampling for general bayesian computation. *Bayesian analysis* 1(4), 833–859 (2006)
44. Theis, L., Oord, A.v.d., Bethge, M.: A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844* (2015)
45. Thin, A., Kotelevskii, N., Denain, J.S., Grinsztajn, L., Durmus, A., Panov, M., Moulines, E.: Metflow: A new efficient method for bridging the gap between Markov chain Monte Carlo and variational inference. *arXiv preprint arXiv:2002.12253* (2020)
46. Thin, A., Kotelevskii, N., Doucet, A., Durmus, A., Moulines, E., Panov, M.: Monte Carlo variational auto-encoders. In: International Conference on Machine Learning. pp. 10247–10257. PMLR (2021)
47. Tom, T.A.L.M.I., Wood, J.T.R.F.: Auto-encoding sequential Monte Carlo. *stat* 1050, 29 (2017)
48. Wolf, C., Karl, M., van der Smagt, P.: Variational inference with Hamiltonian Monte Carlo. *arXiv preprint arXiv:1609.08203* (2016)
49. Wu, H., Köhler, J., Noé, F.: Stochastic normalizing flows. *arXiv preprint arXiv:2002.06707* (2020)
50. Wu, Y., Burda, Y., Salakhutdinov, R., Grosse, R.: On the quantitative analysis of decoder-based generative models. *CoRR arXiv preprint arXiv:1611.04273* (2016)
51. Yu, Y., Zhang, W., Deng, Y.: Frechet inception distance (fid) for evaluating gans
52. Zhang, G., Hsu, K., Li, J., Finn, C., Grosse, R.: Differentiable annealed importance sampling and the perils of gradient noise. *arXiv preprint arXiv:2107.10211* (2021)