
Amortized Monte Carlo Integration

Adam Goliński^{*12} Frank Wood³ Tom Rainforth^{*1}

Abstract

Current approaches to amortizing Bayesian inference focus solely on approximating the posterior distribution. Typically, this approximation is, in turn, used to calculate expectations for one or more target functions—a computational pipeline which is inefficient when the target function(s) are known upfront. In this paper, we address this inefficiency by introducing AMCI, a method for *amortizing Monte Carlo integration* directly. AMCI operates similarly to amortized inference but produces three distinct amortized proposals, each tailored to a different component of the overall expectation calculation. At runtime, samples are produced separately from each amortized proposal, before being combined to an overall estimate of the expectation. We show that while existing approaches are fundamentally limited in the level of accuracy they can achieve, AMCI can theoretically produce arbitrarily small errors for any integrable target function using only a single sample from each proposal at runtime. We further show that it is able to empirically outperform the theoretically optimal self-normalized importance sampler on a number of example problems. Furthermore, AMCI allows not only for amortizing over datasets but also amortizing over target functions.

1. Introduction

At its core, Bayesian modeling is rooted in the calculation of expectations: the eventual aim of modeling is typically to make a decision or to construct predictions for unseen data, both of which take the form of an expectation under the posterior (Robert, 2007). This aim can

^{*}Equal contribution ¹Department of Statistics, University of Oxford, United Kingdom ²Department of Engineering Science, University of Oxford, United Kingdom ³Department of Computer Science, University of British Columbia, Vancouver, Canada. Correspondence to: Adam Goliński <adamg@robots.ox.ac.uk>.

thus be summarized in the form of one or more expectations $\mathbb{E}_{p(x|y)}[f(x)]$, where $f(x)$ is a target function and $p(x|y)$ is the posterior distribution on x for some data y , which we typically only know up to a normalizing constant $p(y)$. More generally, expectations with respect to distributions with unknown normalization constant are ubiquitous throughout the sciences (Robert & Casella, 2013).

Sometimes $f(x)$ is not known up front. Here it is typically convenient to first approximate $p(x|y)$, e.g. in the form of Monte Carlo (MC) samples, and then later use this approximation to calculate estimates, rather than addressing the target expectations directly.

However, it is often the case in practice that a particular target function, or class of target functions, is known a priori. For example, in decision-based settings $f(x)$ takes the form of a loss function, while any posterior predictive distribution constitutes a set of expectations with respect to the posterior, parameterized by the new input. Though often overlooked, it is well established that in such *target-aware* settings the aforementioned pipeline of first approximating $p(x|y)$ and then using this as a basis for calculating $\mathbb{E}_{p(x|y)}[f(x)]$ is suboptimal as it ignores relevant information in $f(x)$ (Hesterberg, 1988; Wolpert, 1991; Oh & Berger, 1992; Evans & Swartz, 1995; Meng & Wong, 1996; Chen & Shao, 1997; Gelman & Meng, 1998; Lacoste-Julien et al., 2011; Owen, 2013; Rainforth et al., 2018b). As we will later show, the potential gains in such scenarios can be arbitrarily large.

In this paper, we extend these ideas to *amortized* inference settings (Stuhlmüller et al., 2013; Kingma & Welling, 2014; Ritchie et al., 2016; Paige & Wood, 2016; Le et al., 2017; 2018a; Webb et al., 2018), wherein one looks to amortize the cost of inference across different possible datasets by learning an artifact that assists the inference process at runtime for a given dataset. Existing approaches do not operate in a target-aware fashion, such that even if the inference network learns proposals that perfectly match the true posterior for every possible dataset, the resulting estimator is still sub-optimal.

To address this, we introduce AMCI, a framework for performing *Amortized Monte Carlo Integration*. Though still based on learning amortized proposals distributions, AMCI varies from standard amortized inference ap-

proaches in three respects. First, it operates in a target-aware fashion, incorporating information about $f(x)$ into the amortization artifacts. Second, rather than using self-normalization, AMCI employs three distinct proposals for separately estimating $\mathbb{E}_{p(x)} [p(y|x) \max(f(x), 0)]$, $\mathbb{E}_{p(x)} [-p(y|x) \min(f(x), 0)]$, and $\mathbb{E}_{p(x)} [p(y|x)]$, before combining these into an overall estimate. This breakdown allows for arbitrary performance improvements compared to self-normalized importance sampling (SNIS). Finally, to account for cases in which multiple possible target functions may be of interest, AMCI also allows for amortization over parametrized functions $f(x; \theta)$.

Remarkably, AMCI is able to achieve an arbitrarily low error at run-time using only a single sample from each proposal given sufficiently powerful amortization artifacts, contrary to the fundamental limitations on the accuracy of conventional amortization approaches. This ability is based around its novel breakdown of the target expectation into separate components, the subsequent utility of which extends beyond the amortized setting we consider here.

2. Background

2.1. Importance Sampling

Importance Sampling (IS), in its most basic form, is a method for approximating an expectation $\mathbb{E}_{\pi(x)} [f(x)]$ when it is either not possible to sample from $\pi(x)$ directly, or when the simple MC estimate, $\frac{1}{N} \sum_{n=1}^N f(x_n)$ where $x_n \sim \pi(x)$, has problematically high variance (Hesterberg, 1988; Wolpert, 1991). Given a proposal $q(x)$ from which we can sample and for which we can evaluate the data, it forms the following estimate

$$\mu := \mathbb{E}_{\pi(x)} [f(x)] = \int f(x) \frac{\pi(x)}{q(x)} q(x) dx \quad (1)$$

$$\approx \hat{\mu} := \frac{1}{N} \sum_{n=1}^N f(x_n) w_n \quad (2)$$

where $x_n \sim q(x)$ and $w_n := \pi(x_n)/q(x_n)$ is known as the importance weight of sample x_n .

In practice, one often does not have access to the normalized form of $\pi(x)$. For example, in Bayesian inference settings, we typically have $\pi(x) = p(x|y) \propto p(x, y)$. Here we can use our samples to both approximate the normalization constant and the unnormalized integral. Thus if $\pi(x) \propto \gamma(x)$, we have

$$\mathbb{E}_{\pi(x)} [f(x)] = \frac{\int \frac{f(x)\gamma(x)}{q(x)} q(x) dx}{\int \frac{\gamma(x)}{q(x)} q(x) dx} \approx \frac{\sum_{n=1}^N f(x_n) w_n}{\sum_{n=1}^N w_n} \quad (3)$$

where $x_n \sim q(x)$, and $w_n := \gamma(x_n)/q(x_n)$. This approach is known as self-normalized importance sampling (SNIS). Conveniently, we can also construct the SNIS estimate lazily by calculating the empirical measure, i.e. stor-

ing weighted samples,

$$\pi(x) \approx \sum_{n=1}^N w_n \delta_{x_n}(x) / \sum_{n=1}^N w_n \quad (4)$$

and then using this to construct the estimate in (3) when $f(x)$ becomes available. As such, we can also think of SNIS as a method for Bayesian inference as, informally speaking, the empirical measure produced can be thought of as an approximation of the posterior.

For a general unknown target, the optimal proposal, i.e. the proposal which results in estimator having lowest possible variance, is the target distribution $q(x) = \pi(x)$ (see e.g. (Rainforth, 2017, 5.3.2.2)). However, this no longer holds if we have some information about $f(x)$. In this target-aware scenario, the optimal behavior turns out to depend on whether we are self-normalizing or not.

For the non-self-normalized case, the optimal proposal can be shown to be $q^*(x) \propto \pi(x) |f(x)|$ (Owen, 2013). Interestingly, in the case where $f(x) \geq 0 \forall x$, this leads to an exact estimator, i.e. $\hat{\mu} = \mu$ (with $\hat{\mu}$ as per (2)). To see this, notice that the normalizing constant for $q^*(x)$ is $\int \pi(x) f(x) dx = \mu$ and hence $q^*(x) = \pi(x) f(x) / \mu$. Therefore, even when $N = 1$, any possible value of the resulting sample x_1 yields an $\hat{\mu}$ satisfying $\hat{\mu} = f(x_1) \pi(x_1) / q^*(x_1) = \mu$.

In the self-normalized case, the optimal proposal instead transpires to be $q^*(x) \propto \pi(x) |f(x) - \mu|$ (Hesterberg, 1988). In this case, one can no longer achieve a zero variance estimator for finite N and nonconstant $f(x)$. Instead, the achievable error is lower bounded by (Owen, 2013)

$$\mathbb{E}[(\hat{\mu} - \mu)^2] \geq \frac{1}{N} \left(\mathbb{E}_{\pi(x)} [|f(x) - \mu|] \right)^2, \quad (5)$$

creating a fundamental limit on the performance of SNIS, even when information about $f(x)$ is incorporated.

Given that these optimal proposals make use of the true expectation μ , we will clearly not have access to them in practice. However, they provide a guide for the desirable properties of a proposal and can be used as targets for adaptive IS methods (see (Bugallo et al., 2017) for a recent review).

2.2. Inference Amortization

Inference amortization involves learning an *amortization artifact* that takes in datasets and produces proposals tailored to the corresponding inference problems. This amortization artifact typically takes the form of a parametrized proposal, $q(x; \varphi(y; \eta))$, which takes in data y and produces proposal parameters using an *inference network* $\varphi(y; \eta)$, which itself has parameters η . When clear from the context, we will use the shorthand $q(x; y, \eta)$ for this proposal.

Though the exact process varies with context, the inference network is usually trained either by drawing latent-data sample pairs from the joint $p(x, y)$ (Paige & Wood, 2016;

Le et al., 2017; 2018b), or by drawing mini-batches from a large dataset using stochastic variational inference approaches (Hoffman et al., 2013; Kingma & Welling, 2014; Rezende et al., 2014; Ritchie et al., 2016). Once trained, it provides an efficient means of approximately sampling from the posterior of a particular dataset, e.g. using SNIS.

Out of several variants, we focus on the method introduced by Paige & Wood (2016), as this is the one AMCI builds upon. In their approach, η is trained to minimize the expectation of $D_{KL} [p(x|y) || q(x; y, \eta)]$ across possible datasets y , giving the objective

$$\begin{aligned} \mathcal{J}(\eta) &= \mathbb{E}_{p(y)} \left[D_{KL} [p(x|y) || q(x; y, \eta)] \right] \\ &= \mathbb{E}_{p(x,y)} [-\log q(x; y, \eta)] + \text{const wrt } \eta \end{aligned} \quad (6)$$

We note that the distribution $p(y)$ over which we are taking the expectation is actually chosen somewhat arbitrarily: it simply dictates how much the network prioritizes a good amortization for one dataset over another; different choices are equally valid and imply different loss functions.

This objective requires us to be able to sample from the joint distribution $p(x, y)$ and it can be optimized using gradient methods since the gradient can be easily evaluated:

$$\nabla_{\eta} \mathcal{J}(\eta) = \mathbb{E}_{p(x,y)} [-\nabla_{\eta} \log q(x; y, \eta)]. \quad (7)$$

3. AMCI

Amortized Monte Carlo integration (AMCI) is a framework for amortizing the cost of calculating expectations $\mu(y, \theta) := \mathbb{E}_{\pi(x;y)} [f(x; \theta)]$. Here y represents changeable aspects of the reference distribution $\pi(x; y)$ (e.g. the dataset) and θ represents changeable parameters of the target function $f(x; \theta)$. The reference distribution is typically known only up to a normalization constant, i.e. $\pi(x; y) = \gamma(x; y)/Z$ where $\gamma(x; y)$ can be evaluated pointwise, but Z is unknown. AMCI can still be useful in settings where Z is known, but here we can simply use its known value rather than constructing a separate estimator.

Amortization can be performed across y and/or θ . When amortizing over y , the function does not need to be explicitly parameterized; we just need to be able to evaluate it pointwise. Similarly, when amortizing over θ , the reference distribution can be fixed. In fact, AMCI can be used for a parameterized set of conventional integration problems $\int_{x \in \mathcal{X}} f(x; \theta) dx$ by exploiting the fact that

$$\int_{x \in \mathcal{X}} f(x; \theta) dx = \mathbb{E}_{\pi(x)} [f(x; \theta) / \pi(x)] \quad (8)$$

for any $\pi(x)$ where $\pi(x) \neq 0 \forall x \in \mathcal{X}$ for which $f(x) \neq 0$.

For consistency of notation with the amortized inference literature, we will presume a Bayesian setting in the rest of this section, i.e. $\pi(x; y) = p(x|y)$ and $\gamma(x; y) = p(x, y)$.

3.1. Estimator

Existing amortized inference methods implicitly evaluate expectations using SNIS (or some other form of self-normalized estimator (Paige & Wood, 2016; Le et al., 2018a)), targeting the posterior as the optimal proposal $q^*(x; y) \approx p(x|y)$. Not only is this proposal suboptimal when information about the target function is available, there is a lower bound on the accuracy the SNIS approach itself can achieve as shown in (5).

AMCI overcomes these limitations by breaking down the overall expectation into separate components and constructing separate estimates for each. We can first break down the target expectation into the ratio of the “unnormlized expectation” and the normalization constant:

$$\begin{aligned} \mu(y, \theta) &:= \mathbb{E}_{p(x|y)} [f(x; \theta)] = \frac{\mathbb{E}_{p(x|y)} [f(x; \theta) p(y)]}{\mathbb{E}_{p(x)} [p(y|x)]} \\ &= \frac{\mathbb{E}_{q_1(x;y,\theta)} \left[\frac{f(x;\theta)p(x,y)}{q_1(x;y,\theta)} \right]}{\mathbb{E}_{q_2(x;y)} \left[\frac{p(x,y)}{q_2(x;y)} \right]} =: \frac{E_1}{E_2} \end{aligned} \quad (9)$$

where $q_1(x; y, \theta)$ and $q_2(x; y)$ are two separate proposals, used respectively for each of the two expectations E_1 and E_2 . We note that the proposal $q_1(x; y, \theta)$ may depend not only on the observed dataset y , but also on the parameters of the target function θ .

We can now generate separate MC estimates for E_1 and E_2 , and take their ratio to estimate the overall expectation:

$$\begin{aligned} \mu(y, \theta) &\approx \hat{\mu}(y, \theta) := \hat{E}_1 / \hat{E}_2 \quad \text{where} \\ \hat{E}_1 &:= \frac{1}{N} \sum_{n=1}^N \frac{f(x'_n; \theta) p(x'_n, y)}{q_1(x'_n; y, \theta)} \quad x'_n \sim q_1(x; y, \theta) \\ \hat{E}_2 &:= \frac{1}{M} \sum_{m=1}^M \frac{p(x_m, y)}{q_2(x_m; y)} \quad x_m \sim q_2(x; y). \end{aligned} \quad (10)$$

The key idea behind AMCI is that we can now **separately train each of these proposals to be good estimators for their respective expectation**, rather than rely on a single proposal to estimate both, as is implicitly the case for SNIS.

Consider, for example, the case where $f(x; \theta) \geq 0$. If $q_1(x; y, \theta) \propto f(x; \theta) p(x|y)$ and $q_2(x; y) \propto p(x|y)$ then both \hat{E}_1 and \hat{E}_2 will form exact estimators (as per Section 2.1), even if $N = M = 1$. Consequently, we achieve an exact estimator for $\mu(y, \theta)$, allowing for arbitrarily large improvements over any SNIS estimator, because SNIS forces $q_1(x; y, \theta)$ and $q_2(x; y)$ to be the same distribution.

More generally, the optimal proposal for E_1 and E_2 are $q_1(x; y, \theta) \propto |f(x; \theta)| p(x|y)$ and $q_2(x; y) \propto p(x|y)$ respectively, with the latter always resulting in an exact estimator for E_2 . Thus the more $|f(x; \theta)| p(x|y)$ varies from $p(x|y)$, the worse the conventional approach of only amortizing

the posterior will perform, while the harder it becomes to construct a reasonable SNIS estimator even when information about $f(x; \theta)$ is incorporated. Separately learning $q_1(x; y, \theta)$ and $q_2(x; y)$ means that each will become a better individual proposal and the overall estimator improves.

It turns out that we do not actually require the previous assumption of $f(x; \theta) \geq 0 \forall x, \theta$ to achieve a zero variance estimator. Specifically, if we let¹

$$f^+(x; \theta) = \max(f(x; \theta), 0) \quad \text{and} \quad (11)$$

$$f^-(x; \theta) = -\min(f(x; \theta), 0) \quad (12)$$

denote truncations of the target function into its positive and negative components (as per the concept of positivisation (Owen, 2013, 9.13)), then we can break down the overall expectation as follows

$$\begin{aligned} & \mu(y, \theta) \\ &= \frac{\mathbb{E}_{q_1^+(x; y, \theta)} \left[\frac{f^+(x; \theta) p(x, y)}{q_1^+(x; y, \theta)} \right] - \mathbb{E}_{q_1^-(x; y, \theta)} \left[\frac{f^-(x; \theta) p(x, y)}{q_1^-(x; y, \theta)} \right]}{\mathbb{E}_{q_2(x; y)} \left[\frac{p(x, y)}{q_2(x; y)} \right]} \\ &=: \frac{E_1^+ - E_1^-}{E_2} \end{aligned} \quad (13)$$

where we now have three expectations and three proposals. Analogously to (10), we can construct estimates for each expectation separately and then combine them:

$$\begin{aligned} \mu(y, \theta) &\approx \hat{\mu}(y, \theta) := (\hat{E}_1^+ - \hat{E}_1^-) / \hat{E}_2 \quad \text{where} \\ \hat{E}_1^+ &:= \frac{1}{N} \sum_{n=1}^N \frac{f^+(x_n^+; \theta) p(x_n^+, y)}{q_1^+(x_n^+; y, \theta)} \quad x_n^+ \sim q_1^+(x; y, \theta) \\ \hat{E}_1^- &:= \frac{1}{K} \sum_{k=1}^K \frac{f^-(x_k^-; \theta) p(x_k^-, y)}{q_1^-(x_k^-; y, \theta)} \quad x_k^- \sim q_1^-(x; y, \theta) \\ \hat{E}_2 &:= \frac{1}{M} \sum_{m=1}^M \frac{p(x_m, y)}{q_2(x_m; y)} \quad x_m \sim q_2(x; y), \end{aligned} \quad (14)$$

which forms the AMCI estimator. The theoretical capability of this estimator is summarized in the following result, the proof for which is given in Appendix B.

Theorem 1. *If the following hold for a given θ and y ,*

$$\mathbb{E}_{p(x)} [f^+(x; \theta) p(y|x)] < \infty \quad (15)$$

$$\mathbb{E}_{p(x)} [f^-(x; \theta) p(y|x)] < \infty \quad (16)$$

$$\mathbb{E}_{p(x)} [p(y|x)] < \infty \quad (17)$$

and we use the corresponding set of optimal proposals $q_1^+(x; y, \theta) \propto f^+(x; \theta) p(x, y)$, $q_1^-(x; y, \theta) \propto f^-(x; \theta) p(x, y)$, and $q_2(x; y) \propto p(x, y)$, then the AMCI

¹Practically, it may sometimes be beneficial to truncate the proposal about another point, c , by instead using $f^+(x; \theta) = \max(f(x; \theta) - c, 0)$ and $f^-(x; \theta) = -\min(f(x; \theta) - c, 0)$, then adding c onto our final estimate.

estimator defined in (14) satisfies

$$\mathbb{E} [\hat{\mu}(y, \theta)] = \mu(y, \theta), \quad \text{Var} [\hat{\mu}(y, \theta)] = 0 \quad (18)$$

for any $N \geq 1$, $K \geq 1$, and $M \geq 1$, such that it forms an exact estimator for that θ, y pair.

Though our primary motivation for developing the AMCI estimator is its attractive properties in an amortization setting, we note that it may still be of use in static expectation calculation settings. Namely, the fact that it can achieve an arbitrarily low mean squared error for a given number of samples means it forms an attractive alternative to SNIS more generally, particularly when we are well-placed to hand-craft highly effective proposals and in adaptive importance sampling settings.

We note that individual elements of this estimator have previously appeared in the literature. For example, the general concept of using multiple proposals has been established in the context of multiple importance sampling (Veach & Guibas, 1995). The use of two separate proposals for the unnormalized target and the normalizing constant (i.e. (10)), on the other hand, was recently independently suggested by Lamberti et al. (2018) in a non-amortized setting. However, we believe that the complete form of the AMCI estimator in (14) has not previously been suggested, nor its theoretical benefits or amortization considered.

3.2. Amortization

To evaluate (14), we need to learn three amortized proposals $q_1^+(x; y, \theta)$, $q_1^-(x; y, \theta)$, and $q_2(x; y)$.

Learning $q_2(x; y)$ is equivalent to the standard inference amortization problem and so we will just use the objective given by (6), as described in section 2.2.

The approaches for learning $q_1^+(x; y, \theta)$ and $q_1^-(x; y, \theta)$ are equivalent, other than the function that is used in the estimators. Therefore, for simplicity, we introduce our amortization procedure in the case where $f(x; \theta) \geq 0 \forall x, \theta$, such that we can need only learn a single proposal, $q_1(x; y, \theta)$, for the numerator as per (10). This trivially extends to the full AMCI setup by separately repeating the same training procedure for $q_1^+(x; y, \theta)$ and $q_1^-(x; y, \theta)$.

3.2.1. FIXED FUNCTION $f(x)$

We first consider the scenario where $f(x)$ is fixed (i.e. we are not amortizing over function parameters θ) and hence in this section we drop the dependence of q_1 on θ .

To learn the parameters η for the first amortized proposal $q_1(x; y, \eta)$, we need to adjust the target in (6) to incorporate the effect of the target function. Let $E_1(y) := \mathbb{E}_{p(x)} [f(x) p(y|x)]$ and $g(x|y) := \frac{f(x) p(x, y)}{E_1(y)}$, i.e. the normalized optimal proposal for q_1 . Naively adjusting (6)

leads to a double intractable objective

$$\begin{aligned} \mathcal{J}'_1(\eta) &= \mathbb{E}_{p(y)} [D_{KL}(g(x|y) || q_1(x; y, \eta))] \\ &= \mathbb{E}_{p(y)} \left[- \int_{\mathcal{X}} \frac{f(x) p(x, y)}{E_1(y)} \log q_1(x; y, \eta) dx \right] \\ &\quad + \text{const wrt } \eta. \end{aligned} \quad (19)$$

Here the double intractability comes from the fact we do not know $E_1(y)$ and, at least at the beginning of the training process, we cannot estimate it efficiently either.

To address this, we use our previous observation that the expectation over $p(y)$ in the above objective is chosen somewhat arbitrarily. Namely, it dictates the relative priority of different datasets y during training and not the optimal proposal for each individual datapoint; disregarding the finite capacity of the network, the global optimum is still always $D_{KL}[g(x|y) || q_1(x; y, \eta)] = 0, \forall y$. We thus maintain a well-defined objective if we choose a different reference distribution over datasets. In particular, if we take the expectation with respect to $h(y) \propto p(y)E_1(y)$, we get

$$\begin{aligned} \mathcal{J}_1(\eta) &= \mathbb{E}_{h(y)} [D_{KL}(g(x|y) || q_1(x; y, \eta))] \\ &= c^{-1} \mathbb{E}_{p(x, y)} [-f(x) \log q_1(x; y, \eta)] \\ &\quad + \text{const wrt } \eta \end{aligned} \quad (20)$$

where $c = \mathbb{E}_{p(y)} [E_1(y)] > 0$ is a positive constant that does not affect the optimization—it is the normalization constant for the distribution $h(y)$ —and can thus be ignored. Each term in this expectation can now be evaluated directly, meaning we can again run stochastic gradient descent algorithms to optimize it. Note that this does not require evaluation of the density $p(x, y)$, only the ability to draw samples.

Interestingly, this choice of $h(y)$ can be interpreted as giving larger importance to the values of y which yield larger $E_1(y)$. Informally, we could think about this choice as attempting to minimizing the L1 errors of our estimates, that is $\mathbb{E}_{p(y)} [|E_1(y) - \hat{E}_1(y)|]$, presuming that the error in our estimation scales as the magnitude of the true value $E_1(y)$.

More generally, if we choose $h(y) \propto p(y)E_1(y)\lambda(y)$ for some positive evaluable function $\lambda : \mathcal{Y} \rightarrow \mathbb{R}^+$, we get a tractable objective of the form

$$\mathcal{J}_1(\eta; \lambda) = \mathbb{E}_{p(x, y)} \left[- \frac{f(x)}{\lambda(y)} \log q_1(x; y, \eta) \right]$$

up to a constant scaling factor and offset. We can thus use this trick to adjust the relative preference given to different datasets, while ensuring the objective is tractable.

3.2.2. PARAMETERIZED FUNCTION $f(x; \theta)$

As previously mentioned, AMCI also allows for amortization over parametrized functions, to account for cases in which multiple possible target functions may be of interest. We can incorporate this by using *pseudo prior* $p(\theta)$ to

generate example parameters during our training.

Analogously to $h(y)$, the choice of $p(\theta)$ determines how much importance we assign to different possible functions that we would like to amortize over. Since, in practice, perfect performance is unattainable over the entire space of θ , the choice of $p(\theta)$ is important and it will have an important effect on the performance of the system.

Incorporating $p(\theta)$ is straightforward: we take the expectation of the fixed target function training objective over θ . In this setting, our inference network φ needs to take θ as input when determining the parameters of q_1 and hence we let $q_1(x; y, \theta, \eta) := q_1(x; \varphi(y, \theta; \eta))$. If $E_1(y, \theta) := \mathbb{E}_{p(x)} [f(x; \theta)p(y|x)]$, $g(x|y; \theta) := f(x; \theta)p(x, y)/E_1(y, \theta)$, and $h(y, \theta) \propto p(y)p(\theta)E_1(y; \theta)$, we get an objective which is analogous to (20):

$$\begin{aligned} \mathcal{J}_1(\eta) &= \mathbb{E}_{h(y, \theta)} [D_{KL}(g(x|y; \theta) || q_1(x; y, \theta, \eta))] \\ &= c^{-1} \cdot \mathbb{E}_{p(x, y)p(\theta)} [-f(x; \theta) \log q_1(x; y, \theta, \eta)] \\ &\quad + \text{const wrt } \eta \end{aligned} \quad (21)$$

where $c = \mathbb{E}_{p(y)p(\theta)} [E_1(y, \theta)] > 0$ is again a positive constant that does not affect the optimization.

3.3. Efficient Training

If $f(x; \theta)$ and $p(x)p(\theta)$ are mismatched, i.e. $f(x; \theta)$ is large in regions where $p(x)p(\theta)$ is low, training by naïvely sampling from $p(x)p(\theta)$ can be inefficient. Instead, it is preferable to try and sample from $g(\theta, x) \propto p(x)p(\theta)f(x; \theta)$. Though this is itself an intractable distribution, it represents a standard, rather than an amortized, inference problem and so it is much more manageable than the overall training. Namely, as the samples do not depend on the proposal we are learning or the datasets, we can carry out this inference process as a pre-training step that is substantially less costly than the problem of training the inference networks itself.

One approach is to construct an MCMC sampler targeting $g(\theta, x)$ to generate the samples, which can be done upfront before training. Another is to use an importance sampler

$$\begin{aligned} \mathcal{J}_1(\eta) &= \text{const wrt } \eta \\ &\quad + c^{-1} \mathbb{E}_{q'(\theta, x)p(y|x)} \left[- \frac{p(\theta)p(x)f(x; \theta)}{q'(\theta, x)} \log q_1(x; y, \theta, \eta) \right] \end{aligned} \quad (22)$$

where $q'(\theta, x)$ is a proposal as close to $g(\theta, x)$ as possible.

In the case of non-parameterized functions $f(x)$, there is no need to take an expectation over $p(\theta)$, and we instead desire to sample from $g(x) \propto p(x)f(x)$.

4. Experiments

Even though AMCI is theoretically able to achieve exact estimators with a finite number of samples, this will rarely be the case for practical problems, for which learning per-




Figure 1: Relative mean squared errors (as per (25)) for [left] the one-dimensional and [right] the five-dimensional tail integral example. The solid lines for each estimator indicate the median of $\delta(y, \theta)$ estimated using a common set of 100 samples from $y, \theta \sim p(y)p(\theta)$, with the corresponding $\delta(y, \theta)$ then each separately estimated using 100 samples of the respective $\hat{\delta}(y, \theta)$. The shading instead shows the estimates from replacing $\delta(y, \theta)$ with the 25% and 75% quantiles of $\hat{\delta}(y, \theta)$ for a given y and θ . The median of $\delta(y, \theta)$ is at times outside of this shaded region as $\delta(y, \theta)$ is often dominated by a few large outliers. The dashed line shows the median of $\delta(y, \theta)$ with the $\delta(y, \theta)$ corresponding to the ReMSE optimal SNIS estimator, namely $(\mathbb{E}_{p(x|y)}[|f(x; \theta) - \mu(y, \theta)|])^2/N$ as per (5), which is itself estimated (with only nominal error) using 10^6 samples. We note that the error for SNIS with q_2 proposal is to a large extent flat because there is not a single sample in the estimator for which $f(x; \theta) > 0$, such that they return $\hat{\mu}(y, \theta) = 0$ and hence give $\delta(y, \theta) = 1$. In Figure (b) the SNIS q_m line reaches the ReMSE value of 10^{18} at $N=2$ and the y-axis limits have been readjusted to allow clear comparison at higher N . This effect is caused by the bias of SNIS: these extremely high errors for SNIS q_m arise when all N samples happen to be drawn from distribution q_1 , for further explanation and the full picture see Figure 5 in Appendix A.

fect proposals is not typically realistic, particularly in amortized contexts (Cremer et al., 2018). It is therefore necessary to test its empirical performance to assert that gains are possible with inexact proposals. To this end, we investigate AMCI’s performance on two illustrative examples.

Our primary baseline is the SNIS approach implicitly used by most existing inference amortization methods, namely the SNIS estimator with proposal $q_2(x; y)$. Though this effectively represents the previous state-of-the-art in amortized expectation calculation, it turns out to be a very weak baseline. We, therefore, introduce another simple approach one could hypothetically consider using: training separate proposals as per AMCI, but then using this to form a mixture distribution proposal for an SNIS estimator. For example, in the scenario where $f(x; \theta) \geq 0 \forall x, \theta$ (such that we only need to learn two proposals), we can use

$$q_m(x; y, \theta) = \frac{1}{2}q_1(x; y, \theta) + \frac{1}{2}q_2(x; y) \quad (23)$$

as an SNIS proposal that takes into account the needs of both E_1 and E_2 . We refer to this method as the *mixture* SNIS estimator and emphasize that it represents a novel amortization approach in its own right.

We also compare AMCI to the theoretically optimal SNIS estimator, i.e. the error bound given by (5). As we will show, AMCI is often able to empirically outperform this bound, thereby giving better performance than *any* approach based on SNIS, whether that approach is amortized

or not. This is an important result and, it particular, it highlights that the potential significance of the AMCI estimator extends beyond the amortized setting we consider here.

We further consider using SNIS with proposal $q_1(x; y, \theta)$. However, this transpires to perform extremely poorly throughout (far worse than $q_2(x; y)$) and so we omit its results from the main paper, giving them in Appendix A.

In all experiments, we use the same number of sample from each proposal to form the estimate (i.e. $N = M = K$).

An implementation for AMCI and our experiments is available at <http://github.com/talesa/amci>.

4.1. Tail Integral Calculation

We start with the conceptually simple problem of calculating tail integrals for Gaussian distributions, namely

$$p(x) = \mathcal{N}(x; 0, \Sigma_1) \quad p(y|x) = \mathcal{N}(y; x, \Sigma_2) \quad (24)$$

$$f(x; \theta) = \prod_{i=1}^D \mathbb{1}_{x_i > \theta_i} \quad p(\theta) = \text{UNIFORM}(\theta; [0, u_D]^D)$$

where D is the dimensionality, we set $\Sigma_2 = I$, and Σ_1 is a fixed covariance matrix (for details see Appendix C).

This problem was chosen because it permits easy calculation of the ground truth expectations by exploiting analytic simplifications, while remaining numerically challenging for values of θ far away from the mean when we do not use these simplifications. We performed one and

five-dimensional variants of the experiment.

We use normalizing flows (Rezende & Mohamed, 2015) to construct our proposals, providing a flexible and powerful means of representing the target distributions. Details are given in Appendix C. Training was done by using importance sampling to generate the values of θ and x as per (22) with $q'(\theta, x) = p(\theta) \cdot \text{HALFNORMAL}(x; \theta, \text{diag}(\Sigma_2))$.

To evaluate AMCI and our baselines we use the relative mean squared error (ReMSE) $\delta(y, \theta) = \mathbb{E}[\hat{\delta}(y, \theta)]$, where

$$\hat{\delta}(y, \theta) = \frac{(\mu(y, \theta) - \hat{\mu}(y, \theta))^2}{\mu(y, \theta)^2} \quad (25)$$

and $\hat{\mu}(y, \theta)$ is our estimate for $\mu(y, \theta)$. We then consider summary statistics across different $\{y, \theta\}$, such as its median when $y, \theta \sim p(y)p(\theta)$.² In calculating this, $\delta(y, \theta)$ was separately estimated for each value of y and θ using 100 samples of $\hat{\delta}(y, \theta)$ (i.e. 100 realizations of the estimator).

As shown in Figure 1, AMCI outperformed SNIS in both the one- and five-dimensional cases. For the one-dimensional example, AMCI significantly outperformed all of SNIS q_2 , SNIS q_m , and the theoretically optimal SNIS estimator. SNIS q_2 , the approach implicitly taken by existing inference amortization methods, typically failed to place even a single sample in the tail of the distribution, even for large N . Interestingly, SNIS q_m closely matched the theoretical SNIS bound, suggesting that this amortized proposal is very close to the theoretically optimal one. However, this still constituted significantly worse performance than AMCI—taking about 10^3 more samples to achieve the same relative error—demonstrating the ability of AMCI to outperform the best possible SNIS estimator.

For the five-dimensional example, AMCI again significantly outperformed our main baseline SNIS q_2 . Though it still also outperformed SNIS q_m , its advantage was less than in one-dimensional case, and it did not outperform the SNIS theoretical bound. SNIS q_m itself did not match the bound as closely as in the one-dimensional example either, suggesting that the proposals learned were worse than in the one-dimensional case. Further comparisons based on using the mean squared error (instead of ReMSE) are given in Appendix A and show qualitatively similar behavior.

4.2. Planning Cancer Treatment

To demonstrate how AMCI might be used in a more real-world scenario, we now consider an illustrative example relating to cancer diagnostic decisions. Imagine that an oncologist is trying to decide whether to administer a treatment to a cancer patient. Because the treatment is highly invasive, they only want to administer it if there is a realis-

²Variability in $\delta(y, \theta)$ between different instances of $\{y, \theta\}$ is considered in Figures 7 and 8 in Appendix A.

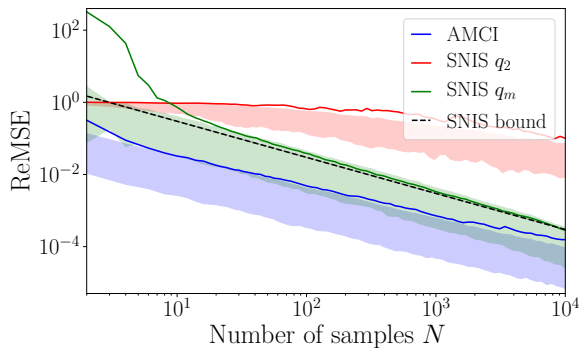


Figure 2: Relative mean squared errors for the cancer example. Conventions as per Figure 1. It is worth noting that it took about 10^4 more samples for the SNIS q_2 estimator to achieve the same level of accuracy as the AMCI estimator.

tic chance of it being successful, i.e. that the tumor shrinks sufficiently to allow a future operation to be carried out. However, they are only able to make noisy observations about the current size of the tumor, and there are various unknown parameters pertaining to its growth, such as the patients predisposition to the treatment. To aid in the oncologists decision, the clinic provides a simulator of tumor evolution, a model of the latent factors required for this simulator, and a loss function for administering the treatment given the final tumor size. We wish to construct an amortization of this simulator, so that we can quickly and directly predict the expected loss function for administering the treatment from a pair of noisy observations of the tumor size taken at separate points in time. A detailed description of the model and proposal setup is in the Appendix C.3.

To evaluate the learned proposals we followed the same procedure as for the tail integral example. Results are presented in Figure 2. AMCI again significantly outperformed the literature baseline of SNIS q_2 —it took about $N = 10^4$ samples for SNIS q_2 to achieve the level of relative error of AMCI for $N = 2$. AMCI further maintained an advantage over SNIS q_m , which itself again closely matched the optimal SNIS estimator. Further comparisons are given in Appendix A and show qualitatively similar behavior.

5. Discussion

In all experiments AMCI performed better than SNIS with either q_2 or q_m for its proposal. Moreover, it is clear that AMCI is indeed able to break the theoretical bound on the achievable performance of SNIS estimators: in some cases AMCI is outperforming the best achievable error by any SNIS estimator, regardless of the proposal the latter uses. Interestingly, the mixture SNIS estimator we also introduced proved to be a strong baseline as it closely matched the theoretical baseline in both experiments. However, such an effective mixture proposal is only possible thanks to learning the multiple inference artifacts we suggest as part of the

AMCI framework, while its performance was still generally inferior to AMCI itself.

We now consider the question of when we expect AMCI to work particularly well compared to SNIS, and the scenarios where it is less beneficial, or potentially even harmful. We first note that scaling with increasing dimensionality is a challenge for both because the importance sampling upon which they rely suffers from the curse of dimensionality. However, the scaling of AMCI should be no worse than existing amortization approaches as each of the amortized proposals is trained in isolation and corresponds to a conventional inference amortization.

We can gain more insights into the relative performance of the two approaches in different settings using an informal asymptotic analysis in the limit of a large number of samples. Assuming $f(x; \theta) \geq 0 \forall x, \theta$ for simplicity,³ then both AMCI and SNIS can be expressed in the form of (10), where for SNIS we set $q_1(x; y, \theta) = q_2(x; y)$, $N = M$, and share samples between the estimators. Separately applying the central limit theorem to \hat{E}_1 and \hat{E}_2 yields

$$\hat{\mu}(y, \theta) = \frac{\hat{E}_1}{\hat{E}_2} \rightarrow \frac{E_1 + \sigma_1 \xi_1}{E_2 + \sigma_2 \xi_2}, \quad \text{as } N, M \rightarrow \infty \quad (26)$$

where $\xi_1, \xi_2 \sim \mathcal{N}(0, 1)$ and

$$\sigma_1 := \frac{1}{N} \text{Var}_{q_1(x; y, \theta)} \left[\frac{f(x; \theta)p(x, y)}{q_1(x; y, \theta)} \right], \quad (27)$$

$$\sigma_2 := \frac{1}{M} \text{Var}_{q_2(x; y)} \left[\frac{p(x, y)}{q_2(x; y)} \right]. \quad (28)$$

Asymptotically, the mean squared error of $\hat{\mu}(y, \theta)$ is dominated by its variance. Thus, by taking a first order Taylor expansion of $\text{Var}[\hat{\mu}(y, \theta)]$ about $1/E_2$, we get, for large M ,

$$\begin{aligned} & \mathbb{E} \left[(\hat{\mu}(y, \theta) - \mu(y, \theta))^2 \right] \\ & \approx \frac{1}{E_2^2} \left(\sigma_1^2 + \sigma_2^2 \mu(y, \theta)^2 - 2\mu(y, \theta) \sigma_1 \sigma_2 \text{Corr}[\xi_1, \xi_2] \right) \\ & = \frac{\sigma_2^2}{E_2^2} \left((\kappa - \text{Corr}[\xi_1, \xi_2])^2 + 1 - \text{Corr}[\xi_1, \xi_2]^2 \right) \quad (29) \end{aligned}$$

where the approximation from the Taylor expansion becomes exact in the limit $M \rightarrow \infty$ and $\kappa := \sigma_1 / (\mu(y, \theta) \sigma_2)$ is a measure of the relative accuracy of the two estimators. See (43) in Appendix D.1 for a more verbose derivation.

For a given value of σ_2 , the value of κ for SNIS is completely dictated by the problem: in general, the larger the mismatch between $f(x; \theta)p(x, y)$ and $p(x, y)$, the larger κ will be. This yields the expected result that the errors for SNIS become large in this setting. For AMCI, we can control κ through ensuring a good proposal for both \hat{E}_1 and \hat{E}_2 , and, if desired, by adjusting M and N (relative to a

³The results trivially generalize to general $f(x)$ with suitable adjustment of the definition of σ_1 .

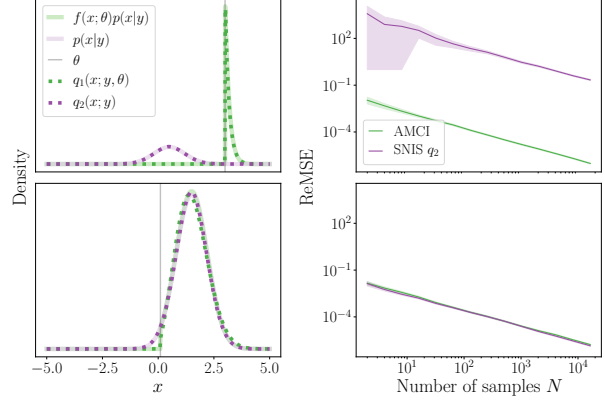


Figure 3: Results for the one-dimensional tail integral model in a setting with large mismatch [top] and low mismatch [bottom], with (y, θ) , respectively $(1, 3)$ and $(3, 0.1)$. The left column illustrates the shape of the proposal q_1 and the achievable quality of fit to $f(x; \theta)p(x|y)$, we see that AMCI is able to learn very accurate proposals in both cases. The right column compares the performance of the AMCI and the SNIS estimators where we see that the gain for AMCI is much larger when the mismatch is large. Uncertainty bands in column two are estimated over a 1000 runs and are almost imperceptibly small.

fixed budget $M + N$). Consequently, we can achieve better errors than SNIS by driving κ down.

On the other hand, as $f(x; \theta)p(x, y)$ and $p(x, y)$ become increasingly well matched, then $\kappa \rightarrow 1$ and we find that AMCI has little to gain over SNIS. In fact, we see that AMCI can potentially be worse than SNIS in this setting: when $f(x; \theta)p(x, y)$ and $p(x, y)$ are closely matched, we also have $\text{Corr}[\xi_1, \xi_2]^2 \approx 1$ for SNIS, such that we observe a canceling effect, potentially leading to very low errors. Achieving $\text{Corr}[\xi_1, \xi_2]^2 \approx 1$ can be more difficult for AMCI, potentially giving rise to a higher error. However, it could be possible to mitigate this by correlating the estimates, e.g. through common random numbers.

To assess if this theory manifests in practice, we revisit our tail integral example, comparing large and small mismatch scenarios. The results, shown in Figure 3, agree with these theoretical findings. In Appendix D we further showing that the reusing of samples for both \hat{E}_1 and \hat{E}_2 in AMCI can be beneficial when the targets are well matched.

More generally, as Theorem 1 tells us that the AMCI estimator can achieve an arbitrarily low error for any given target function, while SNIS cannot, we know that its potential gains are larger the more accurate we are able to make our proposals. As such, as advances elsewhere in the field allow us to produce increasingly effective amortized proposals, e.g. through advanced normalizing flow approaches (Grathwohl et al., 2019; Kingma & Dhariwal, 2018), the larger the potential gains are from using AMCI.

Acknowledgments

We would like to thank Yee Whye Teh for providing helpful discussions at the early stages of the project. AG is supported by the UK EPSRC CDT in Autonomous Intelligent Machines and Systems. FW is supported by DARPA D3M, under Cooperative Agreement FA8750-17-2-0093, Intel under its LBNL NERSC Big Data Center, and an NSERC Discovery grant. TR is supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007–2013)/ERC grant agreement no. 617071. His research leading to these results also received funding from EPSRC under grant EP/P026753/1.

References

- Bugallo, M. F., Elvira, V., Martino, L., Luengo, D., Miguez, J., and Djuric, P. M. Adaptive importance sampling: the past, the present, and the future. *IEEE Signal Processing Magazine*, 34(4):60–79, 2017.
- Chen, M.-H. and Shao, Q.-M. On Monte Carlo methods for estimating ratios of normalizing constants. *The Annals of Statistics*, 25(4):1563–1594, 08 1997.
- Cremer, C., Li, X., and Duvenaud, D. Inference suboptimality in variational autoencoders. *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Enderling, H. and Chaplain, M. A. Mathematical modeling of tumor growth and treatment. *Current pharmaceutical design*, 20–30:4934–40, 2014.
- Evans, M. and Swartz, T. Methods for approximating integrals in statistics with special emphasis on Bayesian integration problems. *Statistical science*, pp. 254–272, 1995.
- Gelman, A. and Meng, X.-L. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185, 05 1998.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. FFFJORD: free-form continuous dynamics for scalable reversible generative models. *International Conference on Learning Representations (ICLR)*, 2019.
- Hahnfeldt, P., Panigrahy, D., Folkman, J., and Hlatky, L. Tumor development under angiogenic signaling. *Cancer Research*, 59(19):4770–4775, 1999.
- Hesterberg, T. C. *Advances in importance sampling*. PhD thesis, Stanford University, 1988.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. Stochastic variational inference. *Journal of Machine Learning Research (JMLR)*, 2013.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. *International Conference on Learning Representations (ICLR)*, 2014.
- Lacoste-Julien, S., Huszár, F., and Ghahramani, Z. Approximate inference for the loss-calibrated Bayesian. *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Lamberti, R., Petetin, Y., Septier, F., and Desbouvries, F. A double proposal normalized importance sampling estimator. *2018 IEEE Statistical Signal Processing Workshop (SSP)*, pp. 238–242, 2018.
- Le, T. A., Baydin, A. G., and Wood, F. Inference compilation and universal probabilistic programming. *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- Le, T. A., Igl, M., Jin, T., Rainforth, T., and Wood, F. Auto-encoding sequential Monte Carlo. *International Conference on Learning Representations (ICLR)*, 2018a.
- Le, T. A., Kosiorek, A. R., Siddharth, N., Teh, Y. W., and Wood, F. Revisiting reweighted wake-sleep. *arXiv:1805.10469*, 2018b.
- Meng, X.-L. and Wong, W. H. Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, 6:831–860, 1996.
- Oh, M.-S. and Berger, J. O. Adaptive importance sampling in Monte Carlo integration. *Journal of Statistical Computation and Simulation*, 41(3-4):143–168, 1992.
- Owen, A. B. *Monte Carlo theory, methods and examples*. 2013.
- Paige, B. and Wood, F. Inference networks for sequential Monte Carlo in graphical models. *Proceedings of the International Conference on Machine Learning (ICML)*, 2016.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems (NIPS)*, 2017.

- Rainforth, T. *Automating inference, learning, and design using probabilistic programming*. PhD thesis, 2017.
- Rainforth, T., Cornish, R., Yang, H., Warrington, A., and Wood, F. On Nesting Monte Carlo Estimators. *Proceedings of the International Conference on Machine Learning (ICML)*, 2018a.
- Rainforth, T., Zhou, Y., Lu, X., Teh, Y. W., Wood, F., Yang, H., and van de Meent, J.-W. Inference trees: Adaptive inference with exploration. *arXiv preprint arXiv:1806.09550*, 2018b.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *Proceedings of the International Conference on Machine Learning (ICML)*, 2014.
- Ritchie, D., Horsfall, P., and Goodman, N. D. Deep amortized inference for probabilistic programs. *arXiv:1610.05735*, 2016.
- Robert, C. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Springer Science & Business Media, 2007.
- Robert, C. and Casella, G. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- Stuhlmüller, A., Taylor, J., and Goodman, N. Learning stochastic inverses. *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- Veach, E. and Guibas, L. J. Optimally combining sampling techniques for Monte Carlo rendering. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pp. 419–428, 1995.
- Webb, S., Goliński, A., Zinkov, R., Siddharth, N., Rainforth, T., Teh, Y. W., and Wood, F. Faithful inversion of generative models for effective amortized inference. *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- Wolpert, R. L. Monte Carlo integration in Bayesian statistical analysis. *Contemporary Mathematics*, 115:101–116, 1991.