
Continuously tempered Hamiltonian Monte Carlo

Matthew M. Graham
School of Informatics
University of Edinburgh

Amos J. Storkey
School of Informatics
University of Edinburgh

Abstract

Hamiltonian Monte Carlo (HMC) is a powerful Markov chain Monte Carlo (MCMC) method for performing approximate inference in complex probabilistic models of continuous variables. In common with many MCMC methods, however, the standard HMC approach performs poorly in distributions with multiple isolated modes. We present a method for augmenting the Hamiltonian system with an extra continuous temperature control variable which allows the dynamic to bridge between sampling a complex target distribution and a simpler unimodal base distribution. This augmentation both helps improve mixing in multimodal targets and allows the normalisation constant of the target distribution to be estimated. The method is simple to implement within existing HMC code, requiring only a standard leapfrog integrator. We demonstrate experimentally that the method is competitive with annealed importance sampling and simulating tempering methods at sampling from challenging multimodal distributions and estimating their normalising constants.

1 INTRODUCTION

Applications of *Markov chain Monte Carlo* (MCMC) methods to perform inference in challenging statistical physics problems were among the earliest uses of the first modern electronic computers [38]. In the years since, MCMC methods have become a mainstay for performing approximate inference in complex probabilistic models in most fields of computational science. Statistical physics in particular has continued to contribute to key developments of MCMC methodology, such as tempering methods [17, 30, 41] and gradient-based dynamics [11].

The aim in MCMC methods is to construct an aperiodic and irreducible Markov chain which leaves a target distribution invariant, so that samples from the chain can be used to form Monte Carlo estimates of expectations with respect to the target distribution. Several general purpose constructions have been developed for specifying valid Markov transition operators for arbitrary target distributions, including Metropolis–Hastings methods [21, 31], slice sampling [9, 35] and Gibbs sampling [13, 16].

While these methods give a basis for constructing chains that will asymptotically give correct inferences, producing algorithms that will give reasonable results in a practical amount of time is still a major challenge. For the restricted case of target distributions defined by densities that are differentiable functions of a real-valued vector variable, *Hamiltonian Monte Carlo* (HMC)¹ [11, 36] provides a framework for defining efficient Markov chains that exploit the gradient of the target density.

Although HMC is able to efficiently explore contiguous regions of high probability density in a target distribution, as with most MCMC methods using local moves it struggles to move between isolated modes in multimodal target distributions [36]. Tempering methods [17, 30, 41] which augment the state space with an temperature variable are often used to improve exploration in multimodal distributions. As the temperature variable is typically discrete however it cannot be included in HMC updates.

In this paper we present an alternative *continuous tempering* approach which instead augments the state with a continuous temperature variable. This allows use of HMC to jointly update both the temperature and original state, making the method simple to use with existing HMC implementations. The proposed approach both improves exploration of multimodal densities and also provides estimates of the typically unknown normalising constant of the target density, which is often an important inferential quantity in its own right for model comparison [15].

¹Originally termed *Hybrid Monte Carlo* in [11].

2 HAMILTONIAN MONTE CARLO

HMC can be used to perform inference whenever the target distribution of interest is defined on a random vector $\mathbf{x} \in \mathbb{R}^D = \mathcal{X}$, by a density function which is differentiable and has support almost everywhere in \mathcal{X} . Distributions with bounded support can often be mapped to an equivalent unbounded distribution by performing a change of variables via a smooth bijective map [8, 24]. We will follow the common convention of defining the target density in a Boltzmann–Gibbs form in terms of a real-valued *potential energy* function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ and a (typically unknown) normalising constant Z ,

$$p_{\mathbf{x}}(\mathbf{x}) = \frac{1}{Z} \exp(-\phi(\mathbf{x})). \quad (1)$$

The key idea of HMC is to simulate a Hamiltonian dynamic in an extended state space to form long-range proposals for a Metropolis accept-reject step with a high probability of acceptance. The *target vector* \mathbf{x} is augmented with a *momentum vector* $\mathbf{p} \in \mathbb{R}^D$. Typically the momentum is chosen to be independent of \mathbf{x} with marginal $p_{\mathbf{p}}(\mathbf{p}) \propto \exp(-\tau(\mathbf{p}))$, with a joint density

$$p_{\mathbf{x},\mathbf{p}}(\mathbf{x}, \mathbf{p}) = p_{\mathbf{x}}(\mathbf{x}) p_{\mathbf{p}}(\mathbf{p}) \propto \exp(-\phi(\mathbf{x}) - \tau(\mathbf{p})). \quad (2)$$

With analogy to classical dynamics, $\tau(\mathbf{x})$ is referred to as the *kinetic energy* and $h(\mathbf{x}, \mathbf{p}) = \phi(\mathbf{x}) + \tau(\mathbf{p})$ is termed the *Hamiltonian*. By construction, marginalising the joint density over the momenta recovers $p_{\mathbf{x}}(\mathbf{x})$.

Each HMC update involves a discrete-time simulation of the canonical Hamiltonian dynamic

$$\frac{d\mathbf{x}}{dt} = \frac{\partial h^\top}{\partial \mathbf{p}} = \frac{\partial \tau^\top}{\partial \mathbf{p}} \quad \frac{d\mathbf{p}}{dt} = -\frac{\partial h^\top}{\partial \mathbf{x}} = -\frac{\partial \phi^\top}{\partial \mathbf{x}} \quad (3)$$

which conserves the Hamiltonian and is time-reversible and volume-preserving. Through choice of an appropriate symplectic integrator such as the popular leapfrog (Störmer–Verlet) scheme, the simulated discrete time dynamic remains exactly time-reversible and volume-preserving and also approximately conserves the Hamiltonian even over long simulated trajectories [28].

The discrete-time dynamic is used to generate a new proposed state $(\mathbf{x}', \mathbf{p}')$ given the current state (\mathbf{x}, \mathbf{p}) and this proposal accepted or rejected in a Metropolis step with acceptance probability $\min\{1, \exp(h(\mathbf{x}, \mathbf{p}) - h(\mathbf{x}', \mathbf{p}'))\}$. Due to the approximate Hamiltonian conservation the acceptance probability is typically close to one, and so HMC is able to make long-range moves in high-dimensional state-spaces while still maintaining high acceptance rates, a significant improvement over the behaviour typical of simpler Metropolis–Hastings methods in high-dimensions.

The energy conservation property which gives this desirable behaviour also however suggests that standard HMC updates are unlikely to move between isolated modes in a target distribution. The Hamiltonian is approximately conserved over a trajectory therefore $\phi(\mathbf{x}') - \phi(\mathbf{x}) \approx \tau(\mathbf{p}) - \tau(\mathbf{p}')$. Typically a quadratic kinetic energy $\tau(\mathbf{p}) = \mathbf{p}^\top \mathbf{M}^{-1} \mathbf{p} / 2$ is used corresponding to a Gaussian marginal density on the momentum. As this kinetic energy is bounded below by zero, the maximum change in potential energy over a trajectory is approximately equal to the initial kinetic energy.

At equilibrium the momenta will have a Gaussian distribution and so the kinetic energy a χ^2 distribution with mean $D/2$ and variance D [2, 36]. If potential energy barriers significantly larger than $\sim D$ separate regions of the configuration state space the HMC updates are unlikely to move across the barriers meaning impractically long sampling runs will be needed for effective ergodicity.

3 THERMODYNAMIC METHODS

A common approach in MCMC methods for dealing with multimodal target distributions is to introduce a concept of temperature. In statistical mechanics, the Boltzmann distribution on a configuration \mathbf{x} of a mechanical system with energy function ϕ and in thermal equilibrium with a heat bath at temperature T is defined by a probability density $\exp(-\beta\phi(\mathbf{x}))/z(\beta)$ where $\beta = (k_B T)^{-1}$ is the *inverse temperature*, k_B is Boltzmann’s constant and $z(\beta)$ is the *partition function*. At high temperatures ($\beta \rightarrow 0$) the density function becomes increasingly flat across the target state space and, correspondingly, energy barriers between different regions of the state space become lower.

In the above statistical mechanics formulation, if $\mathbf{x} \in \mathbb{R}^D$ the distribution in the limit $\beta \rightarrow 0$ has an improper flat density across the target state space. More usefully from a statistical perspective we can use an inverse temperature variable $\beta \in [0, 1]$ to geometrically bridge between a simple *base distribution* with normalised density $\exp(-\psi(\mathbf{x}))$ at $\beta = 0$ and the target distribution at $\beta = 1$

$$\pi(\mathbf{x} | \beta) = \exp(-\beta\phi(\mathbf{x}) - (1 - \beta)\psi(\mathbf{x})), \quad (4)$$

$$z(\beta) = \int_{\mathcal{X}} \exp(-\beta\phi(\mathbf{x}) - (1 - \beta)\psi(\mathbf{x})) d\mathbf{x}. \quad (5)$$

Several approaches to introducing a system temperature into MCMC methods have been proposed. In *simulated tempering* (ST) [30] an ordered set of inverse temperatures are chosen

$$\{\beta_n\}_{n=0}^N : 0 = \beta_0 < \beta_1 < \dots < \beta_N = 1, \quad (6)$$

and a joint distribution defined with density

$$p_{\mathbf{x},\beta}(\mathbf{x}, \beta_n) \propto \pi(\mathbf{x} | \beta_n) \exp(w_n), \quad (7)$$

where $\{w_n\}_{n=0}^N$ are a set of prior weights associated with each inverse temperature. Alternating MCMC updates of $\beta | \mathbf{x}$ and $\mathbf{x} | \beta$ are performed, with \mathbf{x} samples for which $\beta = 1$ converging in distribution to the target. Updates of the inverse temperature variable can propose moves to a limited set of neighbouring inverse temperatures or sample β independently from its conditional $p_{\beta|\mathbf{x}}(\beta_n | \mathbf{x})$.

The ratio of the marginal densities of $p_{\beta}(1)$ and $p_{\beta}(0)$ can be related to the usually unknown normalising constant Z for the target density by

$$Z = \exp(w_0 - w_N) p_{\beta}(1)/p_{\beta}(0), \quad (8)$$

allowing estimation of Z from a ST chain by computing the ratio of counts of samples with $\beta = 1$ and $\beta = 0$. As an alternative [7] proposes to instead use a ‘*Rao-Blackwellised*’ estimator for Z based on the identity

$$p_{\beta}(\beta_n) = \int_{\mathcal{X}} p_{\beta|\mathbf{x}}(\beta_n | \mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x} \quad (9)$$

which indicates $p_{\beta}(\beta_n)$ can be estimated by averaging the conditional probabilities $p_{\beta|\mathbf{x}}(\beta_n | \mathbf{x})$ across samples of \mathbf{x} from the joint in (7). The authors empirically demonstrate this estimator can give significantly improved estimation accuracy over the simpler count-based estimator.

A problem for ST methods is that as $p_{\beta}(\beta_n) \propto z(\beta_n) \exp(w_n)$ and the partition function can vary across several orders of magnitude, the chain can mix poorly between inverse temperatures. This is often tackled with an iterative approach in which initial pilot runs are used to estimate $p_{\beta}(\beta_n)$ and the weights $\{w_n\}_{n=0}^N$ set so as to try to flatten out this marginal distribution [18].

An alternative is *parallel tempering* (PT) [12, 17, 41], where multiple Markov chains on the target state are run in parallel, with the n^{th} chain having an associated inverse temperature β_n defined as in (6). MCMC updates are performed on each chain which leave a distribution with unnormalised density $\pi(\mathbf{x} | \beta_n)$ invariant. Interleaved with these updates, exchanges of the states of the chains at adjacent inverse temperatures (β_n, β_{n+1}) are proposed in a Metropolis–Hastings step. Parallel tempering sidesteps the issue with mixing between inverse temperatures in ST by keeping the inverse temperatures for the chains fixed, at a cost of a significantly increased state size.

Annealed Importance Sampling (AIS) [34] is another thermodynamic ensemble method, closely related to *Tempered Transitions* [33] and often the ‘go-to’ method for normalising constant estimation in the machine learning literature e.g [39, 43]. In AIS an ordered set of inverse temperatures are defined as in (6) and a corresponding series of transition operators $\{T_n\}_{n=1}^{N-1}$. Each T_n leaves a distribution with unnormalised density $\pi(\mathbf{x} | \beta_n)$ invariant.

Assuming the base distribution corresponding to $\beta_0 = 0$ can be sampled from independently, an independent state $\mathbf{x}^{(0)}$ from this distribution is generated and then the transition operators applied in a fixed sequence $T_1 \dots T_{N-1}$ to generate a chain of intermediate states $\{\mathbf{x}^{(n)}\}_{n=1}^{N-1}$. The product of ratios

$$r = \prod_{n=0}^{N-1} (\pi(\mathbf{x}^{(n)} | \beta_{n+1}) / \pi(\mathbf{x}^{(n)} | \beta_n)) \quad (10)$$

is an importance weight for the final state $\mathbf{x}^{(N-1)}$ with respect to the target distribution (1) and also an unbiased estimator for Z . By simulating multiple independent runs of this process, a set of importance weighted samples can be computed to estimate expectations with respect to the target (1) and the weights r used to unbiasedly estimate Z . Due to Jensen’s inequality the unbiased estimate of Z corresponds to a stochastic lower bound on $\log Z$ [20].

In cases where an exact sample from the target distribution can be generated, running reversed AIS from the target to base distribution allow calculating an unbiased estimate of Z^{-1} and so a stochastic upper bound on $\log Z$ [20]. This combination of stochastically lower bounding $\log Z$ by forward AIS runs and stochastically upper bounding by reverse AIS runs is termed *bidirectional Monte Carlo*.

4 CONTINUOUS TEMPERING

In all three of AIS, PT and ST the choice of the discrete inverse temperature schedule (6) is key to getting the methods to perform well in complex high dimensional distributions [1, 3, 33]. To get reasonable performance it may be necessary to do preliminary pilot runs to guide the number of inverse temperatures and spacing between them to use, adding to the computational burden and difficulty of using these methods in a black-box fashion. A natural question is therefore whether it is possible to use a continuously varying inverse temperature variable.

Path sampling [15] proposes this approach, defining a general *path* as a function parametrised by β which continuously maps between the target density $\exp(-\phi(\mathbf{x}))/Z$ at $\beta = 1$ and a base density $\exp(-\psi(\mathbf{x}))$ at $\beta = 0$, with the geometric bridge in (4) a particular example. A joint target $p_{\mathbf{x},\beta}(\mathbf{x}, \beta) \propto \pi(\mathbf{x} | \beta) \rho(\beta)$ is defined with $\rho(\beta)$ a chosen prior on the inverse temperature variable analogous to the weights $\{w_n\}_{n=0}^N$ in simulated tempering. In [15] it is proposed to construct a Markov chain leaving the joint density invariant by alternating updates of $\mathbf{x} | \beta$ and $\beta | \mathbf{x}$. Samples from the joint system can be used to estimate Z via the *thermodynamic integration* identity

$$\log Z = \int_0^1 \int_{\mathcal{X}} \frac{p_{\mathbf{x},\beta}(\mathbf{x}, \beta)}{p_{\beta}(\beta)} \frac{\partial \log \pi(\mathbf{x} | \beta)}{\partial \beta} d\mathbf{x} d\beta. \quad (11)$$

Adiabatic Monte Carlo [3] also proposes using a continuously varying inverse temperature variable, here specifically in the context of HMC. The original Hamiltonian system (\mathbf{x}, \mathbf{p}) is further augmented with a continuous inverse temperature coordinate $\beta \in [0, 1]$. A *contact Hamiltonian* is defined on the augmented system,

$$h_c(\mathbf{x}, \mathbf{p}, \beta) = \beta\phi(\mathbf{x}) + (1 - \beta)\psi(\mathbf{x}) + \frac{1}{2}\mathbf{p}^\top \mathbf{M}^{-1}\mathbf{p} + \log z(\beta) + h_0 \quad (12)$$

this defining a corresponding *contact Hamiltonian flow*,

$$\frac{d\mathbf{x}}{dt} = \frac{\partial h_c}{\partial \mathbf{p}}, \quad \frac{d\mathbf{p}}{dt} = -\frac{\partial h_c}{\partial \mathbf{x}}, \quad \frac{d\beta}{dt} = h_c - \frac{\partial h_c}{\partial \beta}. \quad (13)$$

The contact Hamiltonian flow restricted to the zero level-set of the contact Hamiltonian (which the initial state can always be arranged to lie in by appropriately choosing the arbitrary constant h_0) generates trajectories which exactly conserve the contact Hamiltonian and extended state space volume element, and correspond to the thermodynamical concept of a reversible adiabatic process.

For a quadratic kinetic energy, $\frac{d\beta}{dt}$ is always non-positive and so forward simulation of the contact Hamiltonian flow generates non-increasing trajectories in β (and backwards simulation generates non-decreasing trajectories in β). In the ideal case this allows the inverse temperature range $[0, 1]$ to be coherently traversed without the random-walk exploration inherent to simulated tempering.

Simulating the contact Hamiltonian flow is non-trivial in practice however: the contact Hamiltonian (12) depends on the log partition function $\log z(\beta)$, the partial derivatives of which require computing expectations with respect to $\pi(\mathbf{x} | \beta)$ which for most problems is intractable to do exactly. Moreover the contact flow can encounter meta-stabilities whereby $\frac{d\beta}{dt}$ becomes zero and the flow halts at an intermediate β meaning the flow no longer defines a bijection between $\beta = 0$ and $\beta = 1$. This can be ameliorated by regular resampling of the momenta however this potentially increases random-walk behaviour.

An alternative *extended Hamiltonian approach* for simulating a system with a continuously varying inverse temperature was proposed recently in the statistical physics literature [19]. The inverse temperature of the system is indirectly set via an auxiliary variable, which we will term a *temperature control variable* $u \in \mathbb{R}$. This control variable is mapped to an interval $[s, 1]$, $0 < s < 1$ via a smooth piecewise defined function $\beta : \mathbb{R} \rightarrow [s, 1]$, with the conditions that for a pair of thresholds (θ_1, θ_2) with $0 < \theta_1 < \theta_2$, $\beta(u) = 1 \forall |u| \leq \theta_1$, $\beta(u) = s \forall |u| \geq \theta_2$ and $s < \beta(u) < 1 \forall \theta_1 < |u| < \theta_2$.

Unlike Adiabatic Monte Carlo, an additional momentum variable \mathbf{v} corresponding to u is also introduced. Although

seemingly a minor difference this simplifies the implementation of the approach significantly as the system retains a symplectic structure and can continue to be viewed within the usual Hamiltonian dynamics framework. An *extended Hamiltonian* is defined on the augmented system

$$\tilde{h}(\mathbf{x}, u, \mathbf{p}, \mathbf{v}) = \beta(u)\phi(\mathbf{x}) + \omega(u) + \frac{1}{2}\mathbf{p}^\top \mathbf{M}^{-1}\mathbf{p} + \frac{v^2}{2m} \quad (14)$$

where ω is a ‘confining potential’ on u and m is the mass (marginal variance) associated with v . This extended Hamiltonian is separable with respect to the extended configuration (\mathbf{x}, u) and extended momentum (\mathbf{p}, \mathbf{v}) and so can be efficiently simulated using a standard leapfrog integrator. In [19] the extended Hamiltonian dynamics are integrated using a Langevin scheme without Metropolis adjustment and shown to improve mixing in several molecular dynamics problems.

Due to the condition $\beta(u) = 1 \forall |u| < \theta_1$, the set of sampled configuration states \mathbf{x} which have associated $|u| < \theta_1$ will (assuming the dynamic is ergodic and Metropolis adjustment were used) asymptotically converge in distribution to the target, and so can be used to estimate expectations without any importance re-weighting. The β function is required to be bounded below by a $s > 0$ in [19] due to the base density being bridged to being an improper uniform density on \mathcal{X} . The partition function $z(\beta) \rightarrow \infty$ as $\beta \rightarrow 0$ in this case, which would imply an infinite density for regions in the extended state space where $\beta(u) = 0$. Even with a non-zero lower bound on β , the large variations in $z(\beta(u))$ across different u values can lead to the dynamic poorly exploring the u dimension.

In [19] this issue is tackled by introducing an adaptive history-dependent biasing potential on u to try to achieve a flat density across a bounded interval $|u| < \theta_2$, using for example metadynamics [25]. The resulting non-Markovian updates bias the invariant distribution of the target state however this can be accounted for either by a re-weighting scheme [5], or using a vanishing adaptation.

5 PROPOSED METHOD

As in [19] we define an extended Hamiltonian on an augmented state (\mathbf{x}, u) with associated momenta (\mathbf{p}, \mathbf{v}) ,

$$\tilde{h}(\mathbf{x}, u, \mathbf{p}, \mathbf{v}) = \beta(u)(\phi(\mathbf{x}) + \log \zeta) + (1 - \beta(u))\psi(\mathbf{x}) - \log \left| \frac{\partial \beta}{\partial u} \right| + \frac{1}{2}\mathbf{p}^\top \mathbf{M}^{-1}\mathbf{p} + \frac{v^2}{2m}. \quad (15)$$

As with the approach of [19], this Hamiltonian is separable and the corresponding dynamic can be efficiently simulated with a leapfrog integrator. The reversible and volume-preserving simulated dynamic can then be used

as a proposal generating mechanism on the joint space $(\mathbf{x}, u, \mathbf{p}, \mathbf{v})$ for a Metropolis–Hastings step as in standard HMC. We will term this approach of running HMC in the extended joint space *joint continuous tempering*.

In contrast to [19] we propose to use a smooth monotonically increasing map $\beta : \mathbb{R} \rightarrow [0, 1]$ as the inverse temperature function, with our default choice in all experiments being the logistic sigmoid $\beta(u) = (1 + \exp(-u))^{-1}$.

As previously ψ is the negative logarithm of a simple *normalised* base density, which (as we will motivate in the next section) we will usually choose to be an approximation to the target density (1). Similarly the $\log \zeta$ term will be chosen to be an approximation to $\log Z$.

We can marginalise out the momenta from the joint distribution defined by this Hamiltonian, giving a joint density

$$p_{\mathbf{x},u}(\mathbf{x}, u) \propto \left| \frac{\partial \beta}{\partial u} \right| \exp \left(-\beta(u)(\phi(\mathbf{x}) + \log \zeta) - (1 - \beta(u))\psi(\mathbf{x}) \right). \quad (16)$$

If we define a variable $\beta = \beta(u)$ and use the change of variables formula for a density, we further have that

$$p_{\mathbf{x},\beta}(\mathbf{x}, \beta) \propto \frac{1}{\zeta^\beta} \exp(-\beta\phi(\mathbf{x}) - (1 - \beta)\psi(\mathbf{x})) \quad (17)$$

The bijectivity between u and β is useful as although if simulating a Hamiltonian dynamic we will generally wish to work with u as it is unbounded, the conditional density on β given \mathbf{x} has a tractable normalised form

$$p_{\beta|\mathbf{x}}(\beta | \mathbf{x}) = \frac{\exp(-\beta\Delta(\mathbf{x}))\Delta(\mathbf{x})}{1 - \exp(-\Delta(\mathbf{x}))}, \quad (18)$$

with $\Delta(\mathbf{x}) = \phi(\mathbf{x}) + \log \zeta - \psi(\mathbf{x})$. This corresponds to an exponential distribution with rate parameter $\Delta(\mathbf{x})$ truncated to $[0, 1]$. As an alternative to the joint updates, another option is therefore to form a Markov chain which leaves (17) invariant by alternating independently sampling β from its conditional given \mathbf{x} and performing a transition which leaves the conditional on \mathbf{x} given β invariant, similar to the suggested approach in [15]. We will term this Gibbs sampling type procedure as *Gibbs continuous tempering*.

Further we can use (18) to write the marginal density p_β

$$p_\beta(\beta) = \mathbb{E} \left[\frac{\exp(-\beta\Delta(\mathbf{x}))\Delta(\mathbf{x})}{1 - \exp(-\Delta(\mathbf{x}))} \right]. \quad (19)$$

By integrating the joint (17) over \mathcal{X} we also have that

$$\begin{aligned} p_\beta(\beta) &= \frac{1}{C\zeta^\beta} \int_{\mathcal{X}} \pi(\mathbf{x} | \beta) d\mathbf{x} = \frac{z(\beta)}{C\zeta^\beta} \\ \Rightarrow p_\beta(0) &= \frac{1}{C}, \quad p_\beta(1) = \frac{Z}{C\zeta}, \end{aligned} \quad (20)$$

for an unknown normalising constant C of the joint, and so for MCMC samples $\{\mathbf{x}^{(s)}, \beta^{(s)}\}_{s=1}^S$ from the joint (17)

$$Z = \frac{p_\beta(1)}{p_\beta(0)} \zeta = \lim_{S \rightarrow \infty} \frac{\sum_{s=1}^S (w_1(\mathbf{x}^{(s)}))}{\sum_{s=1}^S (w_0(\mathbf{x}^{(s)}))} \zeta, \quad (21)$$

with $w_0(\mathbf{x}) = \frac{\Delta(\mathbf{x})}{1 - \exp(-\Delta(\mathbf{x}))}$, $w_1(\mathbf{x}) = \frac{\Delta(\mathbf{x})}{\exp(\Delta(\mathbf{x})) - 1}$.

This can be seen to be a continuous analogue of the *Rao-Blackwellised* estimator (9) used in [7]. Similarly we can calculate consistent estimates of expectations with respect to the target density $p_{\mathbf{x}|\beta}(\mathbf{x} | 1) = \exp(-\phi(\mathbf{x}))/Z$ as importance weighted sums

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}) | \beta = 1] &= \frac{\int_{\mathcal{X}} f(\mathbf{x}) p_{\beta|\mathbf{x}}(1 | \mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}}{\int_{\mathcal{X}} p_{\beta|\mathbf{x}}(1 | \mathbf{x}) p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}} \\ &\approx \frac{\sum_{s=1}^S (w_1(\mathbf{x}^{(s)}) f(\mathbf{x}^{(s)}))}{\sum_{s=1}^S (w_1(\mathbf{x}^{(s)}))}. \end{aligned} \quad (22)$$

We can also estimate expectations with respect to the base density $p_{\mathbf{x}|\beta}(\mathbf{x} | 0) = \exp(-\psi(\mathbf{x}))$ using

$$\mathbb{E}[f(\mathbf{x}) | \beta = 0] \approx \frac{\sum_{s=1}^S (w_0(\mathbf{x}^{(s)}) f(\mathbf{x}^{(s)}))}{\sum_{s=1}^S (w_0(\mathbf{x}^{(s)}))}. \quad (23)$$

Often the base density will have known moments (e.g. mean and covariance of a Gaussian base density) which can be compared to the estimates calculated using (23) to check for convergence problems. Convergence of the estimates to the true moments is not a sufficient condition for convergence of the chain to the target joint density (17) but is necessary.

6 CHOOSING A BASE DENSITY

By applying Hölder’s and Jensen’s inequalities we can bound $p_\beta(\beta)$ (see Appendix A for details)

$$\frac{1}{C} \left(\frac{Z}{\zeta} \right)^\beta \exp(-\beta d^{b \rightarrow t}) \leq p_\beta(\beta) \leq \frac{1}{C} \left(\frac{Z}{\zeta} \right)^\beta, \quad (24)$$

where $d^{b \rightarrow t}$ indicates the *Kullback–Leibler* (KL) divergence from the base to target distribution

$$d^{b \rightarrow t} = \int_{\mathcal{X}} \exp(-\psi(\mathbf{x})) \log \left(\frac{\exp(-\psi(\mathbf{x}))}{\exp(-\phi(\mathbf{x}))/Z} \right) d\mathbf{x}. \quad (25)$$

If $\zeta = Z$ the upper-bound is constant. If additionally we had $d^{b \rightarrow t} = 0$, the bound becomes tight and we would have a flat marginal density on β , which we should improve mixing in the β dimension. In reality we do not know Z and cannot choose a base distribution such that

$d^{b \rightarrow t} = 0$ as we wish to use a simple distribution amenable to exploration. However under the constraint of the base distribution allowing exploration, a reasonable heuristic is to minimise the KL divergence to the target distribution. Further we want ζ as close to Z as possible.

Variational inference is an obvious route for tackling both problems, allowing us to fit a base density in a simple parametric family (e.g. Gaussian) by directly minimising the KL divergence (25) while also giving a lower bound on $\log Z$. In some cases we can use variational methods specifically aimed at the target distribution family. More generally methods such as *Automatic Differentiation Variational Inference* (ADVI) [24] provide a black-box framework for fitting variational approximations to differentiable target densities. In models such as Variational Autoencoders [23, 37] a parametric variational approximation to the target density of interest (e.g. posterior on latent space) is fitted during training of the original model, in which case it provides a natural choice for the base distribution as observed in [43].

A potential problem is that the classes of target distribution that we are particularly interested in applying our approach to — those with multiple isolated modes — are precisely the same distributions that simple variational approximations will tend to fit poorly, the divergence (25) being minimised favouring ‘mode-seeking’ solutions [4], which usually fit only one mode well. This both limits how small the divergence from the base to target can be made, but also crucially is undesirable as we wish to use the base distribution to move between modes in the target. One option would be to instead to minimise the reversed form of the KL divergence from the target to base distribution, this tending to produce ‘mode-covering’ solutions that match global moments (and can also be used to produce an analogous lower bound on the marginal density to that for $d^{b \rightarrow t}$ in (24) as shown in Appendix A). Methods such as *expectation propagation* (EP) [32] do allow moment-matching approximations to be found and may be a good option in some cases. Another possibility would be to use an alternative divergence in a variational framework that favours mode-covering solutions, with methods using the χ -divergence [10] and Rényi divergence [29] having been recently proposed in this context.

A further alternative is to fit multiple local variational approximations $\{q_i(\mathbf{x})\}_{i=1}^L$ by minimising the variational objective from multiple random parameter initialisations (discarding any duplicate solutions as measured by some distance tolerance between the variational parameters), each approximating a single mode well. We can then combine these local approximations into a global approximation $q(\mathbf{x})$, for example using a mixture model

$$q(\mathbf{x}) = \frac{1}{\zeta} \sum_{i=1}^L \exp(\ell_i) q_i(\mathbf{x}), \quad \zeta = \sum_{i=1}^L \exp(\ell_i), \quad (26)$$

with ℓ_i the final variational objective value for q_i ($\log Z$ minus the KL divergence from q_i to target distribution). If the local approximations had non-overlapping support this would lead to a global approximation which is guaranteed to be at least as close in KL divergence as any of the local approximations and a $\log \zeta$ which is at least as tight a lower bound on $\log Z$ as any of the individual ℓ_i [45]. A mixture distribution is unlikely to itself be a good choice of base distribution however, as it will tend to be multimodal. We can therefore instead use a base distribution with moments matched to the fitted mixture, e.g. a Gaussian $\exp[-\psi(\mathbf{x})]$ with mean and covariance matched to the mean and covariance of the mixture $q(\mathbf{x})$.

For complex target distributions it will sometimes remain difficult to find a reasonable ψ and $\log \zeta$ which well approximate ϕ and $\log Z$. If the KL divergence from the base to target distribution is high, the bounds on the marginal in (24) become increasingly loose and this tends to lead to low densities on intermediate inverse temperatures. This reduces the ability of the MCMC dynamic to move between the inverse temperatures corresponding to the target and base distributions and so the mixing gain from the augmentation. Similarly from (20) the density ratio $p_\beta(1)/p_\beta(0)$ is $\exp(\log Z - \log \zeta)$, and so for large $|\log Z - \log \zeta|$ the dynamic will tend to spend most of the time close to $\beta = 1$ if $\log Z > \log \zeta$ or $\beta = 0$ if $\log Z < \log \zeta$. In the former case there will be limited gain from using the extended space while in the latter the variance of the estimator in (22) will be comparable to directly using an importance sampling estimator for the target using samples from the base distribution.

A natural approach to try to ameliorate these effects is to use an iterative ‘bootstrap’ method. An initial ψ and $\log \zeta$ are chosen for example using a Gaussian variational approximation to the target density as described above. A Markov chain which leaves the joint density (17) invariant is then simulated. The samples from this chain can then be used to both compute an estimate for $\log Z$ using (21) and updated estimates of the target density mean and covariance using (22). A new Gaussian base density can then be specified with the updated mean and covariance estimates and $\log \zeta$ chosen to be the updated $\log Z$ estimate. Samples of the new joint density can then be used to update $\log \zeta$ and ψ again and so on. This is analogous to the iterative approaches often used in ST to choose the prior weights on the inverse temperatures [7, 18].

7 EXPERIMENTS

We performed a series of experiments comparing our proposed approaches to existing methods. Code for running all experiments is available at <https://git.io/cthmc>. Appendix A shows an additional illustrative example.

7.1 BOLTZMANN MACHINE RELAXATIONS

As a first experiment we performed inference in Gaussian mixture relaxations of a set of ten synthetic Boltzmann machine distributions [44]. The parameters of the Boltzmann machine distributions were randomly generated so that the corresponding relaxations are highly multimodal and so challenging to explore well. A 2D projection of samples from one generated distribution illustrating this multimodality is shown in Figure 5 in Appendix C.

The moments of the relaxation distributions can be calculated from the moments of the original discrete Boltzmann machine distribution, which for models with a small number of binary units D_B (30 in our experiments) can be computed exactly by exhaustive iteration across the 2^{D_B} discrete states. This allows ground truth moments to be calculated against which convergence can be checked. The parametrisation used is described in Appendix B. A Gaussian base density and approximate normalising constant ζ was fit to each the 10 relaxation target densities by matching moments to a mixture of variational Gaussian approximations (individually fitted using a mean-field approach based on the underlying Boltzmann machine distribution) as described in section 6.

Plots showing the *root mean squared error* (RMSE) in estimates of $\log Z$ and the mean and covariance of the relaxation distribution against computational run time for different sampling methods are shown in Figure 1. The RMSE values are normalised by the RMSEs of the corresponding estimated moments used in the base density (and $\log \zeta$) such that values below unity indicate an improvement in accuracy over the variational approximation. The curves shown are RMSEs averaged over 10 independent runs for each of the 10 generated parameter sets, with the filled regions indicating ± 3 standard errors of the mean. The free parameters of all methods were tuned on one parameter set and these values then fixed across all runs. All methods used a shared Theano [42] implementation running on a Intel Core i5-2400 quad-core CPU for the HMC updates and so run times are roughly comparable.

For ST, *Rao Blackwellised* estimators were used as described in [7], with HMC-based updates of the target state $\mathbf{x} | \beta$ interleaved with independent sampling of $\beta | \mathbf{x}$, and 1000 β_n values used. For AIS, HMC updates were used for the transition operators and separate runs with 1000, 5000 and 10000 β_n values used to obtain estimates at different run times. For the tempering approaches run times correspond to increasing numbers of MCMC samples.

The two *continuous tempering* (CT) approaches, Gibbs CT and joint CT, both dominate in terms of having lower average RMSEs in all three moment estimates across all run times, with joint CT showing marginally better per-

formance on estimates of $\log Z$ and $\mathbb{E}[\mathbf{x}]$. The tempering approaches seem to outperform AIS here, possibly as the highly multimodal nature of the target densities favours the ability of tempered dynamics to move the inverse temperature both up and down and so in and out of modes in the target density, unlike AIS where the fixed temperature updates are more likely to end up with chains confined to a single mode after the initial transitions for low β_n .

7.2 HIERARCHICAL REGRESSION MODEL

As a second experiment, we apply our joint continuous tempering approach to perform Bayesian inference in a hierarchical regression model for predicting indoor radon measurements [14]. To illustrate the ease of integrating our approach in existing HMC-based inference software, this experiment was performed with the Python package PyMC3 [40], with its ADVI feature used to fit the base density and its implementation of the adaptive *No U-Turn Sampler* (NUTS) [22] HMC variant used to sample from the extended space.

The regression target in the dataset is measurements of the amount of radon gas $y^{(i)}$ in $N = 919$ households. Two continuous regressors $\mathbf{x}^{(i)}$ and one categorical regressor $c^{(i)}$ are provided per household. A multilevel regression model defined by the factor graph in 2a was used. The model includes five scalar parameters ($\sigma_{\mathbf{a}}, \mu_{\mathbf{a}}, \sigma_{\mathbf{b}}, \mu_{\mathbf{b}}, \epsilon$), an 85-dimensional intercept vector \mathbf{a} and a two-dimensional regressor coefficients vector \mathbf{b} , giving 92 parameters in total. As an example task, we consider inferring the marginal likelihood of the data under the model. Estimation of the marginal likelihood from MCMC samples of the target density alone is non-trivial, with approaches such as the harmonic mean estimator having high variance. Here we try to establish if our approach can be used in a black-box fashion to compute a reasonable estimate of the marginal likelihood.

As our ‘ground truth’ we use a large batch of long AIS runs (average across 100 runs of 10000 inverse temperatures) on a separate Theano implementation of the model. We use ADVI to fit a diagonal covariance Gaussian variational approximation to the target density and use this as the base density. NUTS chains, initialised at samples from the base density, were then run on the extended space for 2500 iterations. The samples from these chain were used to compute estimates of the normalising constant (marginal likelihood) using the estimator (21). The results are shown in Figure 2. It can be seen that estimates from the NUTS chains in the extended continuously tempered space quickly converge to a marginal likelihood estimate very close to the AIS estimate, and significantly improve over the final lower bound on the marginal likelihood that ADVI converges to.




Figure 1: RMSEs in empirical moments estimated from MCMC samples against run time for various thermodynamic ensemble MCMC methods run on Gaussian Boltzmann machine relaxation target distributions. All RMSEs are relative to the RMSE of the corresponding approximate moments calculated using the moment-matched variational mixtures, so values below 1 represent improvement on deterministic approximation. For AIS points across time axis represent increasing number of inverse temperatures: $(1, 5, 10) \times 10^3$. For ST, Gibbs CT and joint CT curves show RMSEs for expectations calculated with increasing number of samples from chains. All curves / points show mean across 10 runs for each of 10 generated parameter sets. Filled regions / error bars show ± 3 standard errors of mean.

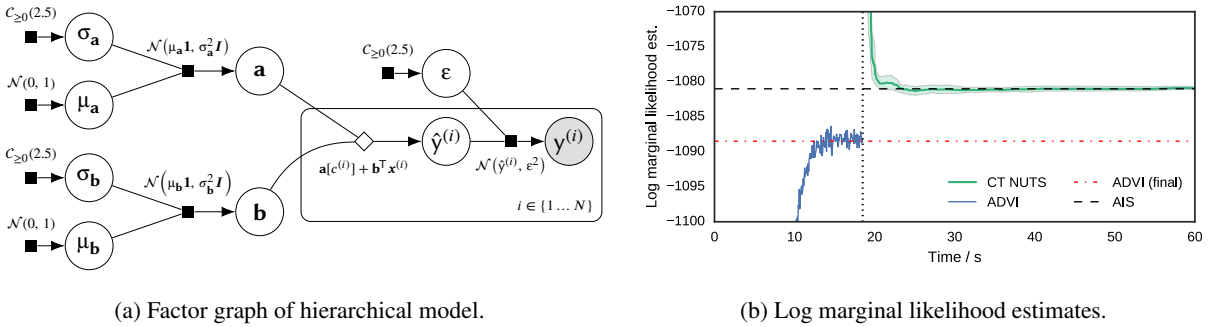



Figure 2: (a) Factor graph for hierarchical regression model for *Radon* data. Factor notation - $\mathcal{N}(\mu, \sigma)$: normal distribution with mean μ , (co-)variance σ ; $C_{\geq 0}(\gamma)$ half-Cauchy distribution with scale γ ; \diamond deterministic factor. (b) Log marginal likelihood estimates against run time for hierarchical regression model. Black dashed line shows estimated log marginal likelihood from a long AIS run which is used as a proxy ground truth. The noisy blue curve shows the *evidence lower bound* ADVI objective over training and the red dot-dashed line the final converged value used for $\log \zeta$. The green curve shows log marginal likelihood estimates using samples from NUTS chains running on the extended joint density in the estimator (21), with the run time corresponding to increasing samples being included in the estimator (offset by initial ADVI run time). Curve shows mean over 10 runs and filled region ± 3 standard errors of mean.

7.3 GENERATIVE IMAGE MODELS


For our final experiments, we compared the efficiency of our CT approaches to ST and AIS for marginal likelihood estimation in decoder-based generative models for images. Use of AIS in this context was recently proposed in [43]. Specifically we estimate the joint marginal likelihood of 1000 generated binary images under the Bernoulli decoder distribution of two *importance weighted autoencoder* (IWAE) [6] models. Each IWAE has one stochastic hidden layer and a 50-dimensional latent space, with the two models trained on binarised versions of the MNIST [27] and Omniglot [26] datasets using the code at <https://github.com/yburda/iwae>. The generated images used are shown in Appendix D.

By performing inference on the per-image posterior densities on the latent representation given image, the joint marginal likelihood of the images can be estimated as the product of estimates of the normalising constants of the individual posterior densities. The use of generated images allows *bidirectional Monte Carlo* (BDMC) [20] to be used to ‘sandwich’ the marginal likelihood with stochastic upper and lower bounds formed with long forward and backward AIS runs (averages over 16 independent runs with 10000 inverse temperatures as used in [43]).

As the per-image latent representations are conditionally independent given the images, chains on all the posterior densities can be run in parallel, with the experiments in this section run on a NVIDIA Tesla K40 GPU to ex-



(a) MNIST log marginal likelihood estimates.



(b) Omniglot log marginal likelihood estimates.

Figure 3: Estimates of the log joint marginal likelihood of 1000 generated images under the Bernoulli decoder distributions of two IWAE models trained on the MNIST and Omniglot datasets against computation time. The black / red dashed lines show stochastic upper / lower bounds calculated using long BDMC runs. For AIS points across time axis red represent increasing number of inverse temperatures: (50, 100, 200, 500, 1000, 2000). For ST, Gibbs CT and joint CT curves show estimates calculated with an increasing number of samples from chains. All curves / points show mean across 10 runs. Filled regions / error bars show ± 3 standard errors of mean.

exploit this inherent parallelism. The encoder of the trained IWAE models is an inference network which outputs the mean and diagonal covariance of a Gaussian variational approximation to the posterior density on the latent representation given an image and so was used to define per-image Gaussian base densities as suggested in [43]. Similarly the per-image $\log \zeta$ values were set using importance-weighted variational lower bound estimates for the per-image marginal likelihoods.

The results are shown in Figure 3, with the curves / points showing average results across 10 independent runs and filled regions / bars ± 3 standard error of means for the estimates. Here Gibbs CT and AIS perform similarly, with joint CT converging less quickly and simulated tempering significantly less efficient. The quick convergence of AIS and Gibbs CT here suggests the posterior densities are relatively easy for the dynamics to explore and well matched by the Gaussian base densities, limiting the gains from any more coherent exploration of the extended space by the joint CT updates. The higher per-leapfrog-step costs of the HMC updates in the extended space therefore mean the joint CT approach is less efficient overall here. The poorer performance of simulated tempering here is in part due to the generation of the discrete random indices becoming a bottleneck in the GPU implementation.

8 DISCUSSION

The approach we have presented is a simple but powerful extension to existing tempering methods which can both help exploration of distributions with multiple isolated modes and allow estimation of the normalisation constant of the target distribution. A key advantage of the joint CT method is its ease of implementation - it simply requires

running HMC in an extended state space and so can easily be used for example within existing probabilistic programming software such as PyMC3 [40] and Stan [8]. By updating the temperature jointly with the original target state, it is also possible to leverage adaptive HMC variants such as NUTS [22] to perform tempered inference in a ‘black-box’ manner without the need to separately tune the updates of the inverse temperature variable.

The Gibbs CT method also provides a relatively black-box framework for tempering. Compared to ST it removes the need to choose the number and spacing of discrete inverse temperatures and also replaces generation of a discrete random variate from a categorical distribution when updating β given \mathbf{x} (which as seen in Section 7.3 can become a computational bottleneck) with generation of a truncated exponential variate (which can be performed efficiently by inverse transform sampling). Compared to the joint CT approach, the Gibbs approach is less simple to integrate in to existing HMC code due to the separate β updates, but eliminates the need to tune the temperature control mass value m and achieved similar or better sampling efficiency in the experiments in Section 7.

Acknowledgements

We thank Ben Leimkuhler for introducing us to the extended Hamiltonian continuous tempering approach and to the anonymous reviewers of an earlier workshop paper for their useful feedback. This work was supported in part by grants EP/F500386/1 & BB/F529254/1 for the University of Edinburgh Doctoral Training Centre in Neuroinformatics & Computational Neuroscience from the UK Engineering & Physical Sciences Research Council (EPSRC), Biotechnology & Biological Sciences Research Council (BBSRC) & Medical Research Council (MRC).

References

- [1] G. Behrens, N. Friel, and M. Hurn. Tuning tempered transitions. *Statistics and computing*, 2012.
- [2] M. Betancourt. A general metric for Riemannian manifold Hamiltonian Monte Carlo. In *Geometric Science of Information*. 2013.
- [3] M. Betancourt. Adiabatic Monte Carlo. *arXiv preprint arXiv:1405.3489*, 2014.
- [4] C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [5] M. Bonomi, A. Barducci, and M. Parrinello. Reconstructing the equilibrium Boltzmann distribution from well-tempered metadynamics. *Journal of computational chemistry*, 2009.
- [6] Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- [7] D. Carlson, P. Stinson, A. Pakman, and L. Paninski. Partition functions from Rao-Blackwellized tempered sampling. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- [8] B. Carpenter, A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 2016.
- [9] P. Damlén, J. Wakefield, and S. Walker. Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 1999.
- [10] A. B. Dieng, D. Tran, R. Ranganath, J. Paisley, and D. M. Blei. The χ -divergence for approximate inference. *arXiv preprint arXiv:1611.00328*, 2016.
- [11] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 1987.
- [12] D. J. Earl and M. W. Deem. Parallel tempering: theory, applications, and new perspectives. *Physical Chemistry Chemical Physics*, 2005.
- [13] A. E. Gelfand and A. F. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 1990.
- [14] A. Gelman and J. Hill. *Data analysis using regression and multilevel/hierarchical models*, chapter 12: Multilevel Linear Models. Cambridge University Press, 2006.
- [15] A. Gelman and X.-L. Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, 1998.
- [16] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 1984.
- [17] C. J. Geyer. Markov chain Monte Carlo maximum likelihood. *Computer Science and Statistics*, 1991.
- [18] C. J. Geyer and E. A. Thompson. Annealing Markov chain Monte Carlo with applications to ancestral inference. *Journal of the American Statistical Association*, 1995.
- [19] G. Gobbo and B. J. Leimkuhler. Extended Hamiltonian approach to continuous tempering. *Physical Review E*, 2015.
- [20] R. Grosse, Z. Ghahramani, and R. P. Adams. Sandwiching the marginal likelihood using bidirectional Monte Carlo. *arXiv preprint arXiv:1511.02543*, 2015.
- [21] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 1970.
- [22] M. D. Hoffman and A. Gelman. The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 2014.
- [23] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- [24] A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. Automatic differentiation variational inference. *arXiv preprint arXiv:1603.00788*, 2016.
- [25] A. Laio and M. Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 2002.
- [26] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 2015.
- [27] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.
- [28] B. Leimkuhler and S. Reich. *Simulating Hamiltonian dynamics*. Cambridge University Press, 2004.
- [29] Y. Li and R. E. Turner. Rényi divergence variational inference. In *Advances in Neural Information Processing Systems*, 2016.
- [30] E. Marinari and G. Parisi. Simulated tempering: a new Monte Carlo scheme. *Europhysics Letters*, 1992.
- [31] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 1953.
- [32] T. P. Minka. Expectation propagation for approximate Bayesian inference. *Proceedings of the Seventeenth conference on Uncertainty in Artificial Intelligence*, 2001.
- [33] R. M. Neal. Sampling from multimodal distributions using tempered transitions. *Statistics and Computing*, 1996.
- [34] R. M. Neal. Annealed importance sampling. *Statistics and Computing*, 2001.
- [35] R. M. Neal. Slice sampling. *Annals of statistics*, 2003.
- [36] R. M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2011.
- [37] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of The 31st International Conference on Machine Learning*, 2014.
- [38] C. Robert and G. Casella. A short history of Markov Chain Monte Carlo: subjective recollections from incomplete data. *Statistical Science*, 2011.
- [39] R. Salakhutdinov and I. Murray. On the quantitative analysis of deep belief networks. In *Proceedings of the 25th International Conference on Machine Learning*, 2008.
- [40] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2016.
- [41] R. H. Swendsen and J.-S. Wang. Replica Monte Carlo simulation of spin-glasses. *Physical Review Letters*, 1986.
- [42] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, 2016.
- [43] Y. Wu, Y. Burda, R. Salakhutdinov, and R. Grosse. On the quantitative analysis of decoder-based generative models. *arXiv preprint arXiv:1611.04273*, 2016.
- [44] Y. Zhang, Z. Ghahramani, A. J. Storkey, and C. A. Sutton. Continuous relaxations for discrete Hamiltonian Monte Carlo. In *Advances in Neural Information Processing Systems*, 2012.
- [45] O. Zoby. Mean field inference for the Dirichlet process mixture model. *Electronic Journal of Statistics*, 2009.