

---

# AMAGOLD: Amortized Metropolis Adjustment for Efficient Stochastic Gradient MCMC

---

Ruqi Zhang  
Cornell University

A. Feder Cooper  
Cornell University

Christopher De Sa  
Cornell University

## Abstract

Stochastic gradient Hamiltonian Monte Carlo (SGHMC) is an efficient method for sampling from continuous distributions. It is a faster alternative to HMC: instead of using the whole dataset at each iteration, SGHMC uses only a subsample. This improves performance, but introduces bias that can cause SGHMC to converge to the wrong distribution. One can prevent this using a step size that decays to zero, but such a step size schedule can drastically slow down convergence. To address this tension, we propose a novel second-order SG-MCMC algorithm—AMAGOLD—that *infrequently* uses Metropolis-Hastings (M-H) corrections to remove bias. The infrequency of corrections amortizes their cost. We prove AMAGOLD converges to the target distribution with a fixed, rather than a diminishing, step size, and that its convergence rate is at most a constant factor slower than a full-batch baseline. We empirically demonstrate AMAGOLD’s effectiveness on synthetic distributions, Bayesian logistic regression, and Bayesian neural networks.

## 1 Introduction

Markov chain Monte Carlo (MCMC) methods play an important role in Bayesian inference. They work by constructing a Markov chain with the desired distribution as its equilibrium distribution; one samples from the chain and, as the algorithm converges to its equilibrium, the samples drawn reflect the desired distribution (Metropolis et al., 1953; Duane et al., 1987; Horowitz, 1991; Neal et al., 2011). Although MCMC is a powerful

technique, when the sampled distribution depends on a very large dataset, its performance is often limited by the dataset’s size—usually by the cost of computing sums over the entire dataset.

One approach for scaling MCMC is to decouple the algorithm from the size of the dataset, using stochastic gradients in lieu of full-batch gradients (Welling and Teh, 2011; Chen et al., 2014; Ding et al., 2014). This family of methods is called Stochastic gradient MCMC (SG-MCMC). For example, stochastic gradient Langevin dynamics (SGLD) replaces the gradient with a stochastic estimate in first order Langevin Monte Carlo (LMC) (Welling and Teh, 2011). The second-order analog of SGLD is stochastic gradient Hamiltonian Monte Carlo (SGHMC). SGHMC can be thought of as a stochastic version of the popular Hamiltonian Monte Carlo (HMC) algorithm and second-order Langevin dynamics (L2MC) (Chen et al., 2014). These algorithms have been particularly useful in Bayesian neural networks, and many variants have been proposed to increase their sampling efficiency and accuracy (Ding et al., 2014; Ahn et al., 2012; Ma et al., 2015; Zhang et al., 2017). Table 1 provides a clarifying summary of the algorithms considered in this paper.

The runtime efficiency benefits of such SG-MCMC methods also come with a drawback: stochastic gradients introduce bias. Bias can cause convergence to a stationary distribution that differs from the one we wanted to sample from, and usually comes from two sources: converting the continuous-time process into discrete gradient updates, and noise from stochastic gradient estimates. These sources of error are in a sense unavoidable because they are also the source of SG-MCMC’s runtime efficiency. Discretization with a large step size  $\epsilon$  (instead of diminishing  $\epsilon \rightarrow 0$ ) allows SG-MCMC to quickly move around its state space, and using stochastic gradients is key for scalability.

The standard approach for removing bias from a Markov chain is to introduce a *Metropolis-Hastings (M-H) correction* (Metropolis et al., 1953). This involves rejecting some fraction of the chain’s transitions

Algorithm	Exact?	Stochastic Gradient?
AMAGOLD	Yes	Yes
L2MC	Yes	No
HMC	Yes	No
SGHMC	No	Yes

Table 1: Comparing 2nd order MCMC methods.

to restore the correct stationary distribution. Naïvely applying M-H to SG-MCMC algorithms is often computationally prohibitive because the M-H step typically needs to sum over the entire dataset. Performing this expensive computation every iteration would defeat the purpose of using stochastic gradients to improve performance. Thus, more sophisticated techniques are needed to achieve efficient, unbiased sampling for SG-MCMC.

In this paper, we show that asymptotic exactness is possible *without being prohibitively expensive*. Specifically, we propose *Amortized Metropolis-Adjusted stochastic Gradient second-Order Langevin Dynamics (AMAGOLD)*. It achieves asymptotic exactness for SGHMC by using an M-H correction step and does so without obliterating the performance gains provided by stochasticity. Our key insight is to apply the M-H step *infrequently*. Rather than computing it every update, AMAGOLD performs it every  $T$  steps ( $T > 0$ ). We prove this is sufficient to remove bias while also improving performance by amortizing the M-H correction cost over  $T$  steps. We develop both reversible and non-reversible AMAGOLD variants and prove both converge to the desired distribution. We also prove a convergence rate relative to using full-batch gradients, which cleanly captures the effect of using stochastic gradients. This result provides insight about the trade-off between minibatching speed-ups and the convergence rate. Our results also show the noise from stochastic gradients has a provably bounded effect on convergence. In summary, our contributions are as follows:

- We introduce AMAGOLD, an efficient, asymptotically-exact SGHMC variant that *infrequently* applies an M-H correction. We give reversible and non-reversible versions.
- We guarantee AMAGOLD converges to the target distribution, and does so without requiring step size  $\epsilon \rightarrow 0$  or precise noise variance estimation.
- We prove a bound on AMAGOLD’s convergence rate with mild assumptions, measured by the spectral gap. This bound is relative to how fast the algorithm *would have* converged if full-batch gradients were used. This is the first such relative convergence bound we are aware of for SG-MCMC.

---

**Algorithm 1** SGHMC
 

---

```

1: given: Energy  $U$ , initial state  $\theta \in \Theta$ , step size  $\epsilon$ ,
   momentum variance  $\sigma^2$ , friction  $\beta$ 
2: loop
3: optionally, resample momentum:
4:  $r \sim \mathcal{N}(0, \sigma^2)$ 
5: initialize position and momentum:
6:  $r_{\frac{1}{2}} \leftarrow r, \theta_0 \leftarrow \theta$ 
7: for  $t = 1$  to  $T$  do
8:   position update:  $\theta_t \leftarrow \theta_{t-1} + \epsilon\sigma^{-2}r_{t-\frac{1}{2}}$ 
9:   sample noise  $\eta_t \sim \mathcal{N}(0, 4\epsilon\beta\sigma^2)$ 
10:  sample random energy component  $\tilde{U}_t$ 
11:  update momentum:
      
$$r_{t+\frac{1}{2}} \leftarrow r_{t-\frac{1}{2}} - \epsilon\nabla\tilde{U}_t(\theta_t) - 2\epsilon\beta r_{t-\frac{1}{2}} + \eta_t$$

12: end for
13: new values:  $(\theta, r) \leftarrow (\theta_T, r_{T+\frac{1}{2}})$ 
14:  $\triangleright$  no M-H step
15: end loop
    
```

---

- We validate our convergence guarantees empirically. Comparing to SGHMC, AMAGOLD is more robust to hyperparameters. Regarding performance, AMAGOLD is competitive with full-batch baselines on synthetic and real-world datasets, and outperforms SGHMC on various tasks.

## 2 Related Work

Our work is situated within a rich literature of SG-MCMC variants that take advantage of stochastic gradient techniques. These methods have demonstrated success on deep neural networks (DNNs) for various tasks (Li et al., 2016c; Gan et al., 2016; Zhang et al., 2020). In particular, second-order SG-MCMC methods like SGHMC, which have a momentum term, have been shown to outperform first-order methods like SGLD on many applications (Chen et al., 2014, 2015). Gao et al. (2018) proves SGHMC’s convergence can be faster than SGLD’s on non-convex problem due to its momentum-based acceleration. SGHMC can also be thought of as a stochastic version of L2MC (Horowitz, 1991) or HMC; we therefore use both L2MC and HMC as experimental full-batch baselines.

Prior work has also studied SGHMC’s convergence properties. Chen et al. (2014) examines its convergence for “asymptotically” small step sizes, in which a continuous-time system governs the dynamics (in contrast, our algorithm is asymptotically exact with a constant step size). Other work proves convergence with high-order integrators (Chen et al., 2015) and obtains non-asymptotic convergence bounds for SGHMC

on non-convex optimization tasks (Gao et al., 2018).

Additional work has studied the properties of first-order M-H adjusted Langevin methods, such as MALA (Grenander and Miller, 1994; Roberts et al., 1996; Roberts and Rosenthal, 1998; Roberts and Stramer, 2002; Stramer and Tweedie, 1999). Dwivedi et al. (2018) derives the mixing time of MALA for strongly log-concave densities, showing it has a better convergence rate than unadjusted Langevin (in comparison, AM-GOLD does not require the assumption of strongly log-concave densities). Korattikara et al. (2014) developed a minibatch M-H approach, which uses subsampling in the M-H correction step, and applied it to correct bias in SGLD. They show cases where SGLD diverges from the target distribution, while SGLD with a minibatched M-H correction performs well.

The work above involves first-order methods. To the best of our knowledge, we are the first to develop an unbiased, efficient second-order SG-MCMC algorithm. We are also the first in this space to use the spectral gap, a traditional metric to evaluate MCMC convergence (Hairer et al., 2014; Levin and Peres, 2017; De Sa et al., 2018). It requires milder assumptions than techniques in prior SG-MCMC work, such as 2-Wasserstein (Raginsky et al., 2017; Dalalyan and Karagulyan, 2019), mean squared error (Vollmer et al., 2016; Chen et al., 2015), and empirical risk (Gao et al., 2018).

### 3 Preliminaries

We start by briefly describing the standard setup of Bayesian inference. Given some dataset  $\mathcal{D}$  and domain  $\Theta$ , suppose we are interested in sampling from the posterior distribution  $\pi(\theta) \propto \exp(-U(\theta))$  where

$$U(\theta) = - \sum_{x \in \mathcal{D}} \log p(x|\theta) - \log p(\theta).$$

$U(\theta)$  is the *energy function*,  $\theta$  ranges over  $\Theta$ , and  $\pi \propto \mu$  denotes  $\pi$  is the unique distribution with PDF proportional to  $\mu$ . One way to compute this distribution is to construct a Markov chain with stationary distribution  $\pi$  and run it to produce a sequence of samples.

A second-order chain, such as HMC, SGHMC, or L2MC (Duane et al., 1987; Horowitz, 1991; Neal et al., 2011), does this by *augmenting* the state space with an additional momentum variable  $r$ , giving joint distribution

$$\pi(\theta, r) \propto \exp(-H(\theta, r)) = \exp\left(-U(\theta) - \frac{1}{2\sigma^2} \|r\|^2\right),$$

where  $H$  is the *Hamiltonian*, which measures the total energy of the system. Note we could replace the norm with any positive definite quadratic form on  $r$ . For

simplicity, we only consider the case of isotropic momentum energy—where the *mass matrix* is  $\sigma^2 I$ . HMC then simulates Hamiltonian dynamics

$$d\theta = \sigma^{-2} r dt, \quad dr = -\nabla U(\theta) dt. \quad (1)$$

The value of the Hamiltonian is preserved under these dynamics, so we must also include transitions that change the value of  $H$  to explore the whole state space. HMC does this by periodically resampling  $r$  from its conditional distribution. L2MC does so by continuously modifying  $r$  with a friction term and added Gaussian noise. Even though (1) preserves  $H$ , the discrete simulation of (1) run by HMC or L2MC *does not* necessarily do so. Therefore both algorithms need an M-H correction step to prevent bias due to discretization.

SGHMC (Algorithm 1) reduces the computational cost of these methods by using a stochastic gradient in lieu of the full-batch gradient  $\nabla U$ . It estimates  $U(\theta)$  using minibatch  $\tilde{\mathcal{D}}$ :

$$\tilde{U}(\theta) \approx -\frac{|\tilde{\mathcal{D}}|}{|\mathcal{D}|} \sum_{x \in \tilde{\mathcal{D}}} \log p(x|\theta) - \log p(\theta).$$

However, using a minibatch introduces noise; naïvely replacing  $U$  by  $\tilde{U}$  leads to divergence from the target distribution. To offset this noise, SGHMC adds the friction term from L2MC (Appendix A). SGHMC uses the leapfrog algorithm to discretize the system (Neal et al., 2011). Notably, SGHMC does not include an M-H correction; to reduce the bias, it requires small  $\epsilon$ .

### 4 Amortized Metropolis Adjustment

*Reversible* Markov chains are a particularly well-studied and well-behaved class of Markov chains. A Markov chain with transition probability operator  $G$  is reversible (also called satisfying the *detailed balance condition*) if for any pair of states  $x$  and  $y$

$$\pi(x)G(x, y) = \pi(y)G(y, x). \quad (2)$$

It is well-known that a chain satisfying (2) has stationary distribution  $\pi$ . An M-H correction constructs a reversible chain  $G$  with stationary distribution  $\pi$  from any Markov chain  $P$  (called the *proposal distribution*) by doing the following at each iteration. First, starting from state  $x$ , sample  $y$  from the proposal distribution  $P(x, y)$ . Second, compute the *acceptance probability*

$$\tau = \min\left(1, \frac{\pi(y)P(y, x)}{\pi(x)P(x, y)}\right).$$

Finally, with probability  $\tau$  transition to state  $y$ ; otherwise, remain in state  $x$ . This correction results in a reversible chain with stationary distribution  $\pi$ ; however, computing  $\tau$  at every step can be costly.

The natural way to amortize the cost of running M-H is to replace the single proposal of baseline M-H with  $T$  proposal-chain steps. This divides its cost among  $T$  iterations of the underlying chain, effectively decreasing it by a factor of  $T$ . For stochastic MCMC, each proposal step can be written as  $P(x, y; \zeta)$ , which denotes the probability of going from state  $x$  to state  $y$  given stochastic sample  $\zeta$  taken from some known distribution. (In minibatched MCMC,  $\zeta$  captures information about which data we sample at that step.) Using this, we can run the following algorithm, starting at  $x$ . First, set  $x_0 = x$ , and run for  $t$  from 0 to  $T - 1$

sample noise  $\zeta_t$ , then sample  $x_{t+1} \sim P(x_t, x_{t+1}; \zeta_t)$ .

Next, set  $y = x_T$ : this is the proposal run an M-H correction on. Finally, compute the acceptance probability

$$\tau = \min \left( 1, \frac{\pi(y)}{\pi(x)} \prod_{t=0}^{T-1} \frac{P(x_{t+1}, x_t; \zeta_t)}{P(x_t, x_{t+1}; \zeta_t)} \right) \quad (3)$$

$$= \min \left( 1, \prod_{t=0}^{T-1} \frac{\pi(x_{t+1})P(x_{t+1}, x_t; \zeta_t)}{\pi(x_t)P(x_t, x_{t+1}; \zeta_t)} \right). \quad (4)$$

and transition to state  $y$  with probability  $\tau$ ; otherwise, remain in state  $x$ .

It is straightforward to see this algorithm results in a reversible chain with stationary distribution  $\pi$ <sup>[1]</sup>. Additionally, this amortized M-H step (3) is easily computed as long as the probabilities  $P(\cdot, \cdot; \zeta)$  are tractable.

We expect this approach will be effective when the M-H step does “not reject too often”. This will certainly be the case when the terms inside the product in (4) all tend to be close to 1, which happens when the proposals  $P(x, y; \eta)$  are “close” to being reversible with stationary distribution  $\pi$ . This is a good heuristic: our amortization approach should be effective when the proposals are close to being reversible.

Unfortunately, this heuristic does not apply to SGHMC’s proposal step since SGHMC and other Hamiltonian-like steps are not close to satisfying the reversibility condition (2). Instead, the natural “reverse” trajectory for a Hamiltonian step reverses the order of the states and *negates the momentum*. The analog of reversibility for this sort of step is *skew-reversibility* (Turitsyn et al., 2011). Given some measure-preserving involution over the state space denoted  $x^\perp$ , a chain  $G$  is skew-reversible if  $\pi(x) = \pi(x^\perp)$  and

$$\pi(x)G(x, y) = \pi(y^\perp)G(y^\perp, x^\perp). \quad (5)$$

Concretely, for Hamiltonian dynamics we use the involution that negates the momentum, i.e.  $(\theta, r)^\perp = (\theta, -r)$ .

It is straightforward to show that a skew-reversible Markov chain also has  $\pi$  as its stationary distribution.<sup>[1]</sup> Such non-reversible chains have attracted a great deal of recent attention because they are more efficient than reversible ones in some situations (Turitsyn et al., 2011; Hukushima and Sakai, 2013; Ma et al., 2016).

A natural consequence of this setup is that we can amortize M-H in the same manner as before, using skew-reversibility in place of reversibility. This gives the same multi-step-proposal algorithm as before, except that the acceptance probability is replaced with

$$\tau = \min \left( 1, \frac{\pi(y^\perp)}{\pi(x)} \prod_{t=0}^{T-1} \frac{P(x_{t+1}^\perp, x_t^\perp; \zeta_t)}{P(x_t, x_{t+1}; \zeta_t)} \right). \quad (6)$$

The resulting corrected chain will be skew-reversible with stationary distribution  $\pi$ .<sup>[1]</sup> Intuitively, this chain will “not reject too often” as long as the proposals  $P$  are “close” to being skew-reversible. Since SGHMC steps are close to being skew-reversible, this is the more natural approach for amortizing M-H, rather than using (3). If one wants to use the well-developed theoretical tools for a reversible chain, it is known that we can recover a reversible chain from a skew-reversible one by simply resampling the momentum at the beginning of the outer loop.<sup>[1]</sup> Note this reversible chain can be different from the one obtained by using condition (3).

## 5 AMAGOLD

We now apply the amortized Metropolis adjustment (AMA) method of Section 4 to second-order SG-MCMC. As a proposal, we use the stochastic leapfrog step that starts in  $(\theta, r)$  and proposes  $(\theta^*, r^*)$  by running

$$\begin{aligned} \theta_0 &= \theta + \frac{1}{2}\epsilon\sigma^{-2}r \\ r^* &= ((1 - \epsilon\beta)r - \epsilon\nabla\tilde{U}_t(\theta_0) + \mathcal{N}(0, 4\epsilon\beta\sigma^2I))/(1 + \epsilon\beta) \\ \theta^* &= \theta_0 + \frac{1}{2}\epsilon\sigma^{-2}r^*. \end{aligned}$$

Applying our amortized M-H correction to this proposal step using the acceptance probability (6) results in AMAGOLD (Algorithm 2). AMAGOLD is, by construction, skew-reversible, and we have the option of making it reversible by resampling the momentum.

AMAGOLD has three key differences compared to SGHMC. First, motivated by the time-reversal-symmetric nature of conditions (2) and (5), we use a clearly time-symmetric update step in the inner loop (compare Line 12 of Algorithm 2 which can be written as  $r_{t+\frac{1}{2}} \leftarrow r_{t-\frac{1}{2}} - \epsilon\nabla\tilde{U}_t(\theta_t) - \epsilon\beta(r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}) + \eta_t$ , with the less clearly symmetric Line 11 of Algorithm 1). Note this is just a different way of writing the algorithm: the update steps could be made equivalent by appropriately setting the hyperparameters. Second,

<sup>1</sup>A detailed proof appears in Appendix B

**Algorithm 2** AMAGOLD

- 
- 1: **given:** Energy  $U$ , initial state  $\theta \in \Theta$ , step size  $\epsilon$ , momentum variance  $\sigma^2$ , friction  $\beta$
  - 2: **loop**
  - 3: **optionally, resample momentum:**  
 $r \sim \mathcal{N}(0, \sigma^2 I)$
  - 4: **initialize momentum, energy acc:**  
 $r_{-\frac{1}{2}} \leftarrow r, \rho_{-\frac{1}{2}} \leftarrow 0$
  - 5: **half position update:**  $\theta_0 \leftarrow \theta + \frac{1}{2}\epsilon\sigma^{-2}r_{-\frac{1}{2}}$
  - 6: **for**  $t = 0$  to  $T - 1$  **do**
  - 7:     **if**  $t \neq 0$  **then**
  - 8:         **position update:**  $\theta_t \leftarrow \theta_{t-1} + \epsilon\sigma^{-2}r_{t-\frac{1}{2}}$
  - 9:     **end if**
  - 10:     **sample noise**  $\eta_t \sim \mathcal{N}(0, 4\epsilon\beta\sigma^2 I)$
  - 11:     **sample random energy component**  $\tilde{U}_t$
  - 12:     **update momentum:**  
 $r_{t+\frac{1}{2}} \leftarrow ((1-\epsilon\beta)r_{t-\frac{1}{2}} - \epsilon\nabla\tilde{U}_t(\theta_t) + \eta_t)/(1+\epsilon\beta)$
  - 13:     **update energy acc:**  
 $\rho_{t+\frac{1}{2}} \leftarrow \rho_{t-\frac{1}{2}} + \frac{1}{2}\epsilon\sigma^{-2}\nabla\tilde{U}_t(\theta_t)^T (r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}})$
  - 14:     **end for**
  - 15:     **half position update:**  
 $\theta_T \leftarrow \theta_{T-1} + \frac{1}{2}\epsilon\sigma^{-2}r_{T-\frac{1}{2}}$
  - 16:     **new values:**  $\theta^* \leftarrow \theta_T, r^* \leftarrow r_{T-\frac{1}{2}}$
  - 17:      $a \leftarrow \exp(U(\theta) - U(\theta^*) + \rho_{T-\frac{1}{2}})$
  - 18:     **with probability**  $\min(1, a)$ ,  
        **update**  $\theta \leftarrow \theta^*, r \leftarrow r^*$  (as long as  $\theta^* \in \Theta$ )
  - 19:     **otherwise update**  $r \leftarrow -r_{-\frac{1}{2}}$
  - 20: **end loop**
- 

we use a type of leapfrog integration that starts and ends the outer loop with a half-position-update (Lines 5 and 15). This too is done in the interest of time-reversal-symmetry. Third, there is an additional term  $\rho$  in AMAGOLD, which we call the *energy accumulator*, which accumulates the log of the product in (6). Computing  $\rho$  requires little extra cost since all its terms are already obtained in the standard update. AMAGOLD is thus unbiased without adding too much cost over SGHMC. The following theorem summarizes AMAGOLD’s asymptotic accuracy. (This follows from the construction; an explicit proof is in Appendix C.)

**Theorem 1.** *Consider the Markov chain described by AMAGOLD (Algorithm 2). If the momentum is resampled (on line 3), then this Markov chain is reversible. Otherwise the Markov chain is skew-reversible. In either case, its stationary distribution is  $\pi$ .*

**Connection to previous methods** AMAGOLD is related to several previous MCMC methods. When

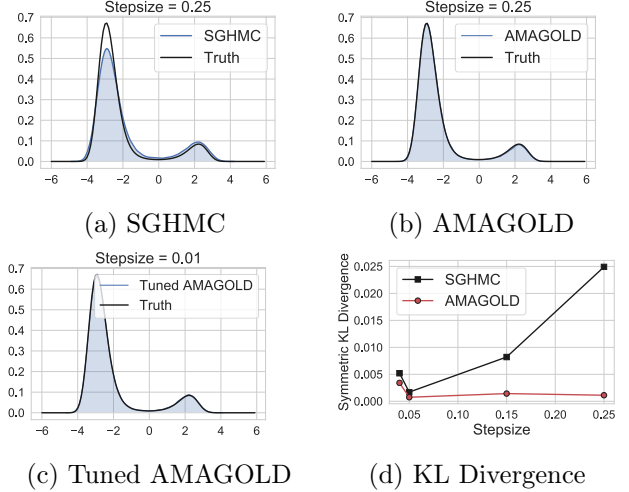


Figure 1: Estimated densities of (a) SGHMC and (b) AMAGOLD for step size 0.25 compared to the ground truth and (c) step size 0.01 for tuned AMAGOLD (see Section 5.2); (d) Comparison of symmetric KL divergence, varying step sizes for SGHMC and AMAGOLD.

using a full-batch gradient, AMAGOLD becomes L2MC with amortized M-H-adjustment. Using a full-batch,  $\beta = 0$ , and resampling, AMAGOLD becomes HMC (Appendix D). If we disable AMAGOLD’s M-H step (and adjust hyperparameters), it becomes SGHMC.

**Illustrating AMAGOLD** To illustrate AMAGOLD is able to achieve unbiased stochastic MCMC, we test our method on a double-well potential (Ding et al., 2014; Li et al., 2016b):

$$U(\theta) = (\theta + 4)(\theta + 1)(\theta - 1)(\theta - 3)/14 + 0.5.$$

The target distribution is proportional to  $\exp(-U(\theta))$ . To simulate stochastic gradients, we let  $\nabla\tilde{U} = \nabla U + \mathcal{N}(0, 1)$ . We show the results of SGHMC and AMAGOLD when  $\beta = 0.25, T = 10$  and  $\epsilon = 0.25$ . Results for  $\epsilon = \{0.05, 0.15\}$  are in Appendix G.1

In Figure 1 and Appendix G.1, the estimated densities of AMAGOLD are very close to the true density on varying step sizes. In contrast, SGHMC does not converge to the correct distribution asymptotically. This is especially the case when the step size is large: SGHMC diverges from the true distribution. These observations validate Theorem 1, as AMA guarantees convergence to the target distribution. To quantitatively measure divergence from the true distribution, we plot the symmetric KL divergence as a function of the step size in Figure 1d. We can see that SGHMC is very sensitive to step size, and may require careful tuning in practice, while AMAGOLD is more robust.

## 5.1 Convergence Rate Analysis

Using stochastic gradients in MCMC can reduce the cost of each iteration. However, this does not mean the overall cost of the algorithm will be less in comparison to its non-stochastic counterpart. Rather, it is possible that the stochastic chain’s convergence rate becomes much slower than the non-stochastic one. To be confident in the effectiveness of an SG-MCMC method, we must rule this out: We must show that the convergence speed of the stochastic chain is not slowed down, or at least not too much, compared to the non-stochastic chain. We do this analysis for AMAGOLD as follows.

Since AMAGOLD can be regarded as stochastic L2MC, we study reversible AMAGOLD’s convergence rate relative to L2MC with an amortized M-H correction. Prior work has used this type of bound to prove the convergence rate of subsampled MCMC methods (De Sa et al., 2018; Zhang and De Sa, 2019). Unlike work that uses 2-Wasserstein, MSE, or empirical risk minimization to evaluate the convergence rate of SG-MCMC, we are the first to use the spectral gap—a traditional metric for evaluating MCMC convergence (Hairer et al., 2014; Levin and Peres, 2017) that is directly related to another common measurement, the mixing time (Levin and Peres, 2017). Our bound only requires mild assumptions compared to prior work (Vollmer et al., 2016; Chen et al., 2015), and we measure convergence to the target distribution directly, rather than empirical risk minimization (Gao et al., 2018).

The spectral gap  $\gamma$  of a reversible Markov chain with transition probability operator  $G$  is defined as the smallest distance between any non-principal eigenvalue of  $G$  and 1, the principal eigenvalue of  $G$  (Levin and Peres, 2017). The spectral gap determines the convergence rate of a Markov chain: a chain with a smaller  $\gamma$  will take longer to converge. To ensure the existence of  $\gamma$ , we assume geometric ergodicity of the full-batch chain (Rudolf, 2011). To bound  $\gamma$ , we assume the covariance of the gradient samples of AMAGOLD is bounded isotropically with

$$\mathbf{E} \left[ (\nabla \tilde{U}(\theta) - \nabla U(\theta))(\nabla \tilde{U}(\theta) - \nabla U(\theta))^T \right] \preceq \frac{V^2}{d} I$$

for some constant  $V > 0$ . This sort of bounded-variance assumption is standard in the analysis of stochastic gradient algorithms.

The following theorem shows that with appropriate hyperparameter settings the convergence rate of AMAGOLD will not be slowed down by more than a constant factor.

**Theorem 2.** *For some parameters  $\epsilon > 0$ ,  $\sigma > 0$ , and  $\beta > 0$ , let  $\bar{\gamma}$  denote the spectral gap of the L2MC chain running with parameters  $(\epsilon, \sigma, \beta)$ . Assume that*

*these parameters are such that  $\epsilon V^2 \leq 4\sigma^2\beta d$ . Define a constant  $c = 1 + \sqrt{\frac{\epsilon V^2}{16\sigma^2\beta T d^2}}$ . Let  $\gamma$  denote the spectral gap of AMAGOLD running with parameters  $(\epsilon, \sigma \cdot c^{-1/4}, \beta \cdot c^{-1/2})$ . Then,*

$$\frac{\gamma}{\bar{\gamma}} \geq \exp \left( -\frac{\epsilon TV^2}{4\sigma^2\beta} - \sqrt{\frac{\epsilon TV^2}{\sigma^2\beta}} \right).$$

This requirement on parameters is easy to satisfy because  $d$  is generally large and  $\epsilon$  is generally small in practice. For the same reason,  $c$  is usually close to 1, so the parameters used by the two chains are very close.

This theorem has three useful takeaways: First, AMAGOLD’s convergence rate is essentially the same as L2MC up to a constant, which will approach 1 as the batch size increases ( $V$  decreases) or  $\epsilon$  decreases. Second, it shows the effect of minibatching on convergence rate: if one reduces the minibatch size (i.e.  $V^2$  increases), they can expect the convergence rate to decrease with a rate of  $\exp(-O(V^2))$ . Third, the theorem outlines a range of parameters (where  $\epsilon TV^2 \ll \sigma^2\beta$ ) over which AMAGOLD converges at a similar rate to the full-batch algorithm.

## 5.2 AMAGOLD in Practice

Here we describe some simple modifications that can further improve AMAGOLD’s performance.

**Minibatch M-H** Amortizing the cost of an M-H correction over  $T$  steps is not always sufficient for achieving good performance on large datasets. This is because calculating the true energy  $U$  requires a scan over the whole dataset. We can further reduce the cost of a single correction by using minibatch M-H to compute the acceptance probability—using a minibatch at line 17 of Algorithm 2. Prior work has estimated the M-H correction using a subset of the data (Korattikara et al., 2014; Bardenet et al., 2014; Maclaurin and Adams, 2015; Seita et al., 2016). These methods are composable with, rather than exclusive with, AMAGOLD and could provide additional speed-ups.

**Tuning the step size** Our experiments on double well potential (Figure 1) show step size significantly influences SGHMC’s performance. Besides being more robust to step size, AMAGOLD’s step size can be more easily tuned. The M-H step’s acceptance probability provides information about whether a step size is desirable. Based on this information, the step size can be tuned automatically to target some fixed acceptance probability during burn-in without affecting convergence. With a fixed step size  $\epsilon = 0.01$ , both AMAGOLD and SGHMC provide poor density estimates due to too small step size. However, when we

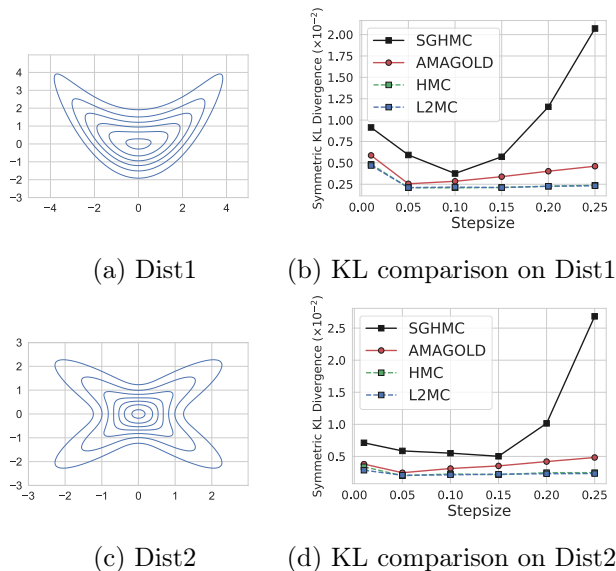


Figure 2: AMAGOLD’s performance against baselines. In (b) and (d) the step size varies from 0.01 to 0.25; the symmetric KL divergence is a function of step size.

let AMAGOLD adjust  $\epsilon$  such that the average M-H acceptance probability is 85%, it estimates the density accurately (Figure 1c, Appendix G.1).

## 6 Experiments

Here we validate our theory empirically and explore the performance of AMAGOLD on a variety of applications. We compare to full-batch baselines HMC and L2MC to show AMAGOLD is more efficient and we compare to SGHMC because, despite exhibiting bias, it is commonly used in the literature. Unless otherwise specified, we use reversible AMAGOLD, meaning we resample the momentum,  $T = 10$  and  $\beta = 0.25$ . We set hyperparameters for AMAGOLD in a similar way as SGHMC (Chen et al., 2014). For simplicity, we do not use the techniques in Section 5.2. Additional details are in Appendix G. The code can be found at <https://github.com/ruqizhang/amagold>.

### 6.1 Synthetic Distributions

We conduct experiments on synthetic two-dimensional distributions (Figures 2a and 2c), which are adapted from (Yin and Zhou, 2018). The analytical expressions are in Appendix G.2.1. We compare our algorithm against three baselines: (1) HMC, (2) L2MC with amortized M-H correction, and (3) SGHMC. HMC and L2MC serve as non-stochastic, unbiased baselines. As in (Chen et al., 2014), we replace  $\nabla U$  by stochastic estimates  $\nabla U = \nabla U + \mathcal{N}(0, I)$  for the stochastic methods. We draw  $5 \times 10^6$  samples and use symmetric




Figure 3: The convergence speed (symmetric KL divergence as a function of iterations) of AMAGOLD compared to L2MC on Dist1 with step size 0.15.

KL divergence as a function of step size to quantitatively evaluate the convergence of the Markov chain. On both distributions, AMAGOLD’s symmetric KL divergence is close to full-batch methods and is much lower than SGHMC’s, especially when the step size is large. This validates our theory that AMAGOLD is unbiased, while SGHMC’s bias increases with step size. See Appendix G.2.3 for more details.

We then verify the theory that AMAGOLD has a comparable convergence rate to L2MC while using stochastic gradient estimates. Specifically, in Figure 3 AMAGOLD’s convergence rate is the same as L2MC’s (up to a constant factor slowdown of about  $10^{-3}$ ). We include runtime comparisons in Appendix G.2.2.

### 6.2 Bayesian Logistic Regression on Real-World Data

We evaluate our method on Bayesian logistic regression using two real-world datasets: *Australian* and *Heart* (Figure 4). We compute the MSE between the estimated and true parameters, obtained from  $10^7$  samples from HMC as in (Li et al., 2016a). AMAGOLD exhibits smaller error than SGHMC on varying step sizes. We show runtime comparisons with step size  $10^{-4}$ . Compared to full-batch HMC and L2MC, AMAGOLD is significantly faster due to minibatching. It is also not much slower than SGHMC, indicating AMA can reduce the cost of adding the M-H step. AMAGOLD’s large error using a large step size is due to a drop in M-H acceptance probability (Appendix G.3). However, this drop can be easily avoided in practice. One can either set the step size such that it achieves a reasonable acceptance rate (usually 20–80%, depending on the application) or use the tuning technique in Section 5.2. With a reasonable acceptance rate, AMAGOLD achieves much lower error compared to SGHMC.

### 6.3 Bayesian Neural Networks

We apply AMAGOLD on Bayesian neural networks. The architecture is a MLP with two-layer with RELU non-linearities. The dataset size is 60000 and we use minibatch size 2000. We use irreversible AMAGOLD

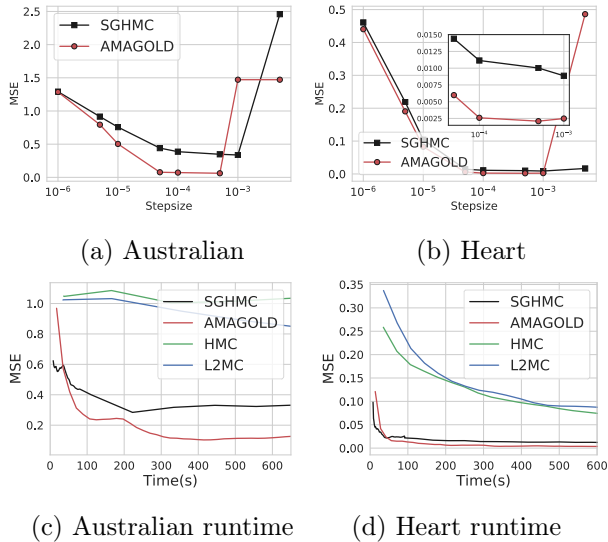


Figure 4: We use two real-world datasets (a) *Australian* (15 covariates, 690 data points) and (b) *Heart* (14 covariates, 270 data points). The minibatch size is 32 and 16, respectively. We collect  $5 \times 10^6$  samples and test step size varying from  $10^{-6}$  to  $5 \times 10^{-3}$ .

Algorithm	$b$	$h = 0.0005$	$h = 0.001$
SGHMC	0.01	$3.69 \pm 0.03$	$3.77 \pm 0.17$
SGHMC	$5e-6$	$89.95 \pm 0.29$	$89.70 \pm 0.91$
AMAGOLD	0.01	$3.63 \pm 0.04$	$3.65 \pm 0.08$
AMAGOLD	$5e-6$	$3.65 \pm 0.10$	$3.63 \pm 0.10$

Table 2: Comparison between AMAGOLD and SGHMC of test error (%)  $\pm$  standard error. We collect 20 samples in total.

since we find it gives better results. Similar to Zhang et al. (2020), to speed up the convergence of the sampling methods, we use SGD with momentum in the first 3 epochs as burn-in and then switch to either SGHMC or AMAGOLD.

**Classification** We evaluate the classification accuracy of AMAGOLD and SGHMC. As in Chen et al. (2014), we reparameterize our algorithm, setting  $v = \epsilon\sigma^{-2}r$ ,  $b = \epsilon\beta$  and  $h = \epsilon^2\sigma^{-2}$  (Appendix F). This equivalent two-parameter reformulated update is similar to SGD with momentum and thus more easily tuned on DNNs. Table 2 shows the test error on various hyperparameter settings. AMAGOLD yields consistent test error, regardless of the hyperparameter values. In contrast, the performance of SGHMC is affected significantly by the hyperparameters. When  $b$  is small, SGHMC diverges. Similar performance of SGHMC has also been reported in Ding et al. (2014).

**Uncertainty Evaluation** We evaluate the sampling performance in terms of uncertainty evaluation, which is important in many ML applications (Lakshminarayanan et al., 2017; Blundell et al., 2015). We test predictive uncertainty estimation on out-of-distribution samples (Lakshminarayanan et al., 2017). The 8 models in Table 2 are tested on the notMNIST dataset (Bulatov, 2011). Since the models have never seen the samples from notMNIST, ideally the predictive distribution should be uniform, which gives the maximum entropy. We plot the empirical CDF for the entropy of the predictive distribution (Figure 5). AMAGOLD provides consistent uncertainty estimations on all settings, which aligns with the classification results. In contrast, when  $b$  is small or  $h$  is large, SGHMC performance suffers; it is overconfident about its prediction.

Both of these experiments indicate that SGHMC is very sensitive to hyperparameters. It needs to be carefully tuned to achieve desired performance on classification and uncertainty estimation. In contrast, AMAGOLD is robust to various hyperparameter settings because it is guaranteed to converge to the target distribution.




Figure 5: Empirical CDF on notMNIST dataset.

## 7 Conclusion

Our work represents a first step toward unbiased, efficient second-order SG-MCMC. We introduced AMAGOLD, which achieves these goals by infrequently applying the computationally-expensive Metropolis-Hasting adjustment step, amortizing the cost across multiple algorithm steps. We prove this is sufficient for convergence to the target distribution, and provide reversible and non-reversible versions for practical use. AMAGOLD’s convergence rate is theoretically guaranteed: the bound captures the trade-off between the speed-up from minibatching and the convergence rate. Lastly, our work is complementary to, rather than exclusive with, other research in stochastic MCMC. In future work it would be interesting to explore combining AMA with other SG-MCMC variants (Ding et al., 2014; Zhang et al., 2017; Ma et al., 2015) and minibatch M-H methods (Korattikara et al., 2014; Bardenet et al., 2014; Maclaurin and Adams, 2015; Seita et al., 2016).



## Acknowledgements

This work was supported in part by Huawei Technologies Co., Ltd.

## References

- Sungjin Ahn, Anoop Korattikara, and Max Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. *arXiv preprint arXiv:1206.6380*, 2012.
- Shigeki Aida. Uniform positivity improving property, Sobolev inequalities, and spectral gaps. *Journal of functional analysis*, 158(1):152–185, 1998.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *International Conference on Machine Learning (ICML)*, pages 405–413, 2014.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- Yaroslav Bulatov. Not MNIST Dataset. 2011. <http://yaroslavvb.blogspot.com/2011/09/notmnist-dataset.html>.
- Changyou Chen, Nan Ding, and Lawrence Carin. On the convergence of stochastic gradient MCMC algorithms with high-order integrators. In *Advances in Neural Information Processing Systems*, pages 2278–2286, 2015.
- Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International conference on machine learning*, pages 1683–1691, 2014.
- Arnak S Dalalyan and Avetik Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 2019.
- Christopher De Sa, Vincent Chen, and Wing Wong. Minibatch gibbs sampling on large graphical models. *arXiv preprint arXiv:1806.06086*, 2018.
- Nan Ding, Youhan Fang, Ryan Babbush, Changyou Chen, Robert D Skeel, and Hartmut Neven. Bayesian sampling using stochastic gradient thermostats. In *Advances in neural information processing systems*, pages 3203–3211, 2014.
- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics letters B*, 195(2):216–222, 1987.
- Raaz Dwivedi, Yuansi Chen, Martin J Wainwright, and Bin Yu. Log-concave sampling: Metropolis-hastings algorithms are fast! *arXiv preprint arXiv:1801.02309*, 2018.
- Masatoshi Fukushima, Yoichi Oshima, and Masayoshi Takeda. *Dirichlet forms and symmetric Markov processes*, volume 19. Walter de Gruyter, 2010.
- Zhe Gan, Chunyuan Li, Changyou Chen, Yunchen Pu, Qinliang Su, and Lawrence Carin. Scalable bayesian learning of recurrent neural networks for language modeling. *arXiv preprint arXiv:1611.08034*, 2016.
- Xuefeng Gao, Mert Gürbüzbalaban, and Lingjiong Zhu. Global convergence of stochastic gradient Hamiltonian Monte Carlo for non-convex stochastic optimization: Non-asymptotic performance bounds and momentum-based acceleration. *arXiv preprint arXiv:1809.04618*, 2018.
- Ulf Grenander and Michael I Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4): 549–581, 1994.
- Martin Hairer, Stuart Andrew M., and Vollmer Sebastian J. Spectral gaps for a Metropolis–Hastings algorithm in infinite dimensions. *The Annals of Applied Probability* 24, no. 6 (2014): 2455–2490, 2014.
- Alan M Horowitz. A generalized guided Monte Carlo algorithm. *Physics Letters B*, 268(2):247–252, 1991.
- K Hukushima and Y Sakai. An irreversible Markov-chain Monte Carlo method with skew detailed balance conditions. In *Journal of Physics: Conference Series*, volume 473, page 012012. IOP Publishing, 2013.
- Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *International Conference on Machine Learning*, pages 181–189, 2014.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.
- David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned stochastic gradient Langevin dynamics for deep neural networks. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016a.

- Chunyuan Li, Changyou Chen, Kai Fan, and Lawrence Carin. High-order stochastic gradient thermostats for Bayesian learning of deep models. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016b.
- Chunyuan Li, Andrew Stevens, Changyou Chen, Yunchen Pu, Zhe Gan, and Lawrence Carin. Learning weight uncertainty with stochastic gradient mcmc for shape classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5666–5675, 2016c.
- Yi-An Ma, Tianqi Chen, and Emily Fox. A complete recipe for stochastic gradient mcmc. In *Advances in Neural Information Processing Systems*, pages 2917–2925, 2015.
- Yi-An Ma, Tianqi Chen, Lei Wu, and Emily B Fox. A unifying framework for devising efficient and irreversible MCMC samplers. *arXiv preprint arXiv:1608.05973*, 2016.
- Dougal Maclaurin and Ryan Prescott Adams. Firefly Monte Carlo: Exact MCMC with subsets of data. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- Radford M Neal et al. MCMC using Hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM J. on Optimization*, 19(4):1574–1609, January 2009. ISSN 1052-6234. doi: 10.1137/070704277. URL <https://doi.org/10.1137/070704277>.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: a nonasymptotic analysis. *arXiv preprint arXiv:1702.03849*, 2017.
- Gareth O Roberts and Jeffrey S Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.
- Gareth O Roberts and Osnat Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and computing in applied probability*, 4(4):337–357, 2002.
- Gareth O Roberts, Richard L Tweedie, et al. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- Daniel Rudolf. Explicit error bounds for markov chain monte carlo. *arXiv preprint arXiv:1108.3201*, 2011.
- Daniel Seita, Xinlei Pan, Haoyu Chen, and John Canny. An efficient minibatch acceptance test for Metropolis-Hastings. *arXiv preprint arXiv:1610.06848*, 2016.
- O Stramer and RL Tweedie. Langevin-type models i: Diffusions with given stationary distributions and their discretizations. *Methodology and Computing in Applied Probability*, 1(3):283–306, 1999.
- Konstantin S Turitsyn, Michael Chertkov, and Marija Vucelja. Irreversible Monte Carlo algorithms for efficient sampling. *Physica D: Nonlinear Phenomena*, 240(4-5):410–414, 2011.
- Sebastian J Vollmer, Konstantinos C Zygalakis, and Yee Whye Teh. Exploration of the (non-) asymptotic bias and variance of stochastic gradient Langevin dynamics. *The Journal of Machine Learning Research*, 17(1):5504–5548, 2016.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011.
- Mingzhang Yin and Mingyuan Zhou. Semi-implicit variational inference. *arXiv preprint arXiv:1805.11183*, 2018.
- Ruqi Zhang and Christopher M De Sa. Poisson-minibatching for gibbs sampling with convergence rate guarantees. In *Advances in Neural Information Processing Systems*, pages 4923–4932, 2019.
- Ruqi Zhang, Chunyuan Li, Jianyi Zhang, Changyou Chen, and Andrew Gordon Wilson. Cyclical stochastic gradient mcmc for bayesian deep learning. *International Conference on Learning Representations*, 2020.
- Yizhe Zhang, Changyou Chen, Zhe Gan, Ricardo Henao, and Lawrence Carin. Stochastic gradient monomial gamma sampler. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*, pages 3996–4005. JMLR. org, 2017.

## Appendix

### A Background of SGHMC

SGHMC, Algorithm 1, reduces the computational cost of full-batch methods by using a stochastic gradient in lieu of  $\nabla U$ . SGHMC estimates  $U(\theta)$  by minibatch  $\tilde{D}$ . However, naïvely replacing  $U$  by  $\tilde{U}$  in HMC will lead to divergence from the target distribution (Chen et al., 2014). Therefore, SGHMC adds the additional friction term of L2MC to offset the noise introduced by using minibatch. That is, SGHMC is simulating the dynamics

$$\begin{aligned} d\theta &= \sigma^{-2} r dt, \\ dr &= -\nabla \tilde{U}(\theta) dt - 2\beta r dt + \mathcal{N}(0, 4(\beta - \hat{\beta})\sigma^2 dt) \end{aligned} \tag{7}$$

where  $\beta$  controls the friction term and  $\hat{\beta}$  is an estimate of the stochastic gradient noise, which is often set to be zero in practice.

### B Proof of Results in Section 4

In this section, we will provide proofs of the results that we asserted in Section 4 about reversibility and skew-reversibility of our algorithms. First, for completeness we re-prove the fact that a skew-reversible chain has stationary distribution  $\pi$ , which is known, but not as well-known as the corresponding result for reversible chains.

**Lemma 1.** *If  $G$  is a skew-reversible chain, that is one that satisfies (5), then  $\pi$  is its stationary distribution.*

*Proof.* Since  $G$  is skew-reversible, by definition it satisfies for any states  $x$  and  $y$  the conditions that  $\pi(x) = \pi(x^\perp)$  and

$$\pi(x)G(x, y) = \pi(y^\perp)G(y^\perp, x^\perp).$$

By combining these two we can easily get

$$\pi(x)G(x, y) = \pi(y)G(y^\perp, x^\perp).$$

Next, summing up over all  $x$  in the whole state space  $\Omega$ ,

$$\sum_{x \in \Omega} \pi(x)G(x, y) = \sum_{x \in \Omega} \pi(y)G(y^\perp, x^\perp) = \pi(y) \sum_{x \in \Omega} G(y^\perp, x^\perp).$$

Since  $\perp$  denotes an involution, it follows that summing up over  $x$  for all  $x$  in the state space is equal to summing up over all  $x^\perp$ , so

$$\sum_{x \in \Omega} \pi(x)G(x, y) = \pi(y) \sum_{x \in \Omega} G(y^\perp, x) = \pi(y),$$

where the last equality follows from the fact that for any Markov chain, the sum of the probabilities of transitioning into all states is always 1. So, we've shown that

$$\sum_{x \in \Omega} \pi(x)G(x, y) = \pi(y)$$

which can be written in matrix form as  $\pi G = \pi$ ; this proves the lemma. □

**Lemma 2.** *The amortized Metropolis-Hastings procedure described in Section 4 using acceptance probability (3) results in a chain that is reversible with stationary distribution  $\pi$ .*

*Proof.* According to the algorithm described in Section 4, the probability density of transitioning from state  $x$  to state  $y \neq x$  via intermediate states  $x = x_0, x_1, x_2, \dots, x_{T-1}, x_T = y$  is

$$\mathbf{E} \left[ \tau \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right],$$

where the expected value here is taken over the randomness used to select the stochastic samples  $\zeta_t$ . This follows from the law of total expectation. This means that the total probability of transitioning from  $x$  to  $y \neq x$  is just the integral of this over the intermediate states

$$G(x, y) = \int \mathbf{E} \left[ \tau \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right] dx_1 \cdot dx_2 \cdots dx_{T-1},$$

Now substituting in the value of  $\tau$  from (3) gives us

$$\begin{aligned} G(x, y) &= \int \mathbf{E} \left[ \min \left( 1, \frac{\pi(y)}{\pi(x)} \prod_{t=0}^{T-1} \frac{P(x_{t+1}, x_t; \zeta_t)}{P(x_t, x_{t+1}; \zeta_t)} \right) \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right] dx_1 \cdot dx_2 \cdots dx_{T-1} \\ &= \int \mathbf{E} \left[ \min \left( \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t), \frac{\pi(y)}{\pi(x)} \prod_{t=0}^{T-1} P(x_{t+1}, x_t; \zeta_t) \right) \right] dx_1 \cdot dx_2 \cdots dx_{T-1}. \end{aligned}$$

Multiplying both sides by  $\pi(x)$ ,

$$\pi(x)G(x, y) = \int \mathbf{E} \left[ \min \left( \pi(x) \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t), \pi(y) \prod_{t=0}^{T-1} P(x_{t+1}, x_t; \zeta_t) \right) \right] dx_1 \cdot dx_2 \cdots dx_{T-1}.$$

From here, the fact that  $G$  is reversible follows directly from a substitution of  $x_t \mapsto x_{T-t}$  in the integral, combined with the observation that the  $\zeta_t$  are i.i.d. and so exchangeable.  $\square$

**Lemma 3.** *The amortized Metropolis-Hastings procedure for skew-reversible chains described in Section 4 using acceptance probability (6) results in a chain that is skew-reversible with stationary distribution  $\pi$ , as long as  $\pi$  satisfies  $\pi(x) = \pi(x^\perp)$ .*

*Proof.* As above, the probability density of transitioning from state  $x$  to state  $y \neq x$  via intermediate states  $x = x_0, x_1, x_2, \dots, x_{T-1}, x_T = y$  is

$$\mathbf{E} \left[ \tau \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right],$$

where the expected value here is taken over the randomness used to select the stochastic samples  $\zeta_t$ . This follows from the law of total expectation. This means that the total probability of transitioning from  $x$  to  $y \neq x$  is just the integral of this over the intermediate states

$$G(x, y) = \int \mathbf{E} \left[ \tau \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right] dx_1 \cdot dx_2 \cdots dx_{T-1},$$

Now substituting in the value of  $\tau$  from (3) gives us

$$\begin{aligned} G(x, y) &= \int \mathbf{E} \left[ \min \left( 1, \frac{\pi(y)}{\pi(x)} \prod_{t=0}^{T-1} \frac{P(x_{t+1}^\perp, x_t^\perp; \zeta_t)}{P(x_t, x_{t+1}; \zeta_t)} \right) \cdot \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t) \right] dx_1 \cdot dx_2 \cdots dx_{T-1} \\ &= \int \mathbf{E} \left[ \min \left( \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t), \frac{\pi(y)}{\pi(x)} \prod_{t=0}^{T-1} P(x_{t+1}^\perp, x_t^\perp; \zeta_t) \right) \right] dx_1 \cdot dx_2 \cdots dx_{T-1}. \end{aligned}$$

Multiplying both sides by  $\pi(x)$ , and leveraging the fact that  $\pi(x) = \pi(x^\perp)$ ,

$$\pi(x)G(x, y) = \int \mathbf{E} \left[ \min \left( \pi(x) \prod_{t=0}^{T-1} P(x_t, x_{t+1}; \zeta_t), \pi(y^\perp) \prod_{t=0}^{T-1} P(x_{t+1}^\perp, x_t^\perp; \zeta_t) \right) \right] dx_1 \cdot dx_2 \cdots dx_{T-1}.$$

From here, the fact that  $G$  is skew-reversible follows directly from a substitution of  $x_t \mapsto x_{T-t}^\perp$  in the integral (which is a valid substitution without introducing an extra constant term because the involution  $\perp$  is measure-preserving by assumption), combined with the observation that the  $\zeta_t$  are i.i.d. and so exchangeable.  $\square$

**Lemma 4.** *A skew-reversible chain will become reversible by resampling the momentum at the beginning of outer loop.*

*Proof.* Assume the chain starts at  $(\theta, r)$  and ends at  $(\theta^*, r^*)$ . By the skew-detailed balance, we have

$$\pi(\theta, r)G((\theta, r), (\theta^*, r^*)) = \pi(\theta^*, -r^*)G((\theta^*, -r^*), (\theta, -r))$$

Since the momentum is resampled and is independent of  $\theta$ , we can integrate it and describe the chain in terms of  $\theta$

$$\pi(\theta)H(\theta, \theta^*) := \int \pi(\theta)\pi(r)G((\theta, r), (\theta^*, r^*))drdr^*$$

Similarly, we have

$$\pi(\theta^*)H(\theta^*, \theta) := \int \pi(\theta^*)\pi(-r^*)G((\theta^*, -r^*), (\theta, -r))dr^*dr$$

By the skew-detailed balance we know

$$\pi(\theta)H(\theta, \theta^*) = \pi(\theta^*)H(\theta^*, \theta)$$

This proves the lemma. □

## C Proof of Theorem 1

*Proof.* First, we consider resampling momentum in Algorithm 2 and will show that the chain is reversible. We consider the probability of starting from  $\theta$  and going through a particular sequence of  $r_{t+\frac{1}{2}}$  and  $\theta_t$  and arriving at  $(\theta^*, r^*)$ . We have  $G(\theta, \theta^*)$ , which is the transition probability from  $\theta$  to  $\theta^*$ , as the following

$$G(\theta, \theta^*) = \mathbf{E} \int \mathbf{P} \left( \theta_0, \theta_1, \dots, \theta_{T-1}, \theta^* \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \min(1, a(\theta)) d\theta_0 \cdots d\theta_{T-1}$$

where  $\theta = \{\theta_0, \dots, \theta_{T-1}, \theta^*\}$  and the expectation is taken over the stochastic energy function samples  $\tilde{U}_t$ .

Next, we want to derive the probability density in terms of  $r$  and  $\eta = \{\eta_0, \dots, \eta_{T-1}\}$ . This involves a change of variables in the PDF formula. We notice that  $\theta$  is a bijective function of  $r$  and  $\eta$ . By the rule of change of variables, we know that

$$\begin{aligned} & \mathbf{P} \left( \theta_0, \theta_1, \dots, \theta_{T-1}, \theta^* \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \min(1, a(\theta)) \\ &= \mathbf{P} \left( r, \eta_0, \dots, \eta_{T-1} \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \min(1, a(\eta, r)) \det^{-1}(D_{(\eta, r)}(\theta)) \end{aligned}$$

where  $D_{(\eta, r)}(\theta, \theta^*)$  is the Jacobian matrix.

To get this Jacobian matrix, we first apply the chain rule,

$$D_{(\eta, r)}(\theta) = D_r(\theta) \cdot D_{(\eta, r)}r$$

where  $r = \{r, r, \dots, r_{T-\frac{1}{2}}\}$ .

Since the derivative of  $\theta_t$  with respect to any  $r_{s-\frac{1}{2}}$  for  $s > t$  is zero, it follows that  $D_r(\theta)$  will be triangular, and so the determinant is just the product of the diagonal entries. From our formula for the update rule,

$$\begin{aligned} \theta_t &= \theta + \frac{1}{2}\epsilon\sigma^{-2}r_{t-\frac{1}{2}}, \text{ for } t = 0, T \\ \theta_t &= \theta_{t-1} + \epsilon\sigma^{-2}r_{-\frac{1}{2}}, \text{ for } t = 1, \dots, T-1. \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{\partial \theta_0}{\partial r_{t-\frac{1}{2}}} &= \frac{1}{2}\epsilon\sigma^{-2}I_d, \text{ for } t = 0, T \\ \frac{\partial \theta_t}{\partial r_{t-\frac{1}{2}}} &= \epsilon\sigma^{-2}I_d, \text{ for } t = 1, \dots, T-1. \end{aligned}$$

It follows that

$$\det(D_{\mathbf{r}}(\boldsymbol{\theta})) = \frac{1}{4^d} (\epsilon\sigma^{-2})^{(T+1)d}.$$

Similarly, the derivative of  $\eta_t$  with respect to any  $r_{s-\frac{1}{2}}$  for  $s > t$  is zero, it follows that  $D_{(\boldsymbol{\eta}, \mathbf{r})}\mathbf{r}$  will be triangular, and so the determinant is just the product of the diagonal entries. That is,

$$D_{\mathbf{r}}(\boldsymbol{\eta}, r) = \prod_{t=0}^{T-1} \frac{\partial \eta_t}{\partial r_{t+\frac{1}{2}}}.$$

From our original formula for the update rule,

$$r_{t+\frac{1}{2}} = r_{t-\frac{1}{2}} - \epsilon \nabla \tilde{U}_t(\theta_t) - \epsilon\beta \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) + \eta_t$$

we have

$$(1 + \epsilon\beta)r_{t+\frac{1}{2}} = r_{t-\frac{1}{2}} - \epsilon \nabla \tilde{U}_t(\theta_t) - \epsilon\beta r_{t-\frac{1}{2}} + \eta_t,$$

and so

$$\frac{\partial \eta_t}{\partial r_{t+\frac{1}{2}}} = (1 + \epsilon\beta)I_d.$$

It follows that

$$\det(D_{(\boldsymbol{\eta}, \mathbf{r})}\mathbf{r}) = (1 + \epsilon\beta)^{-Td}$$

Now we can get that

$$\begin{aligned} \det(D_{(\boldsymbol{\eta}, \mathbf{r})}(\boldsymbol{\theta})) &= \det(D_{\mathbf{r}}(\boldsymbol{\theta})) \cdot \det(D_{(\boldsymbol{\eta}, \mathbf{r})}\mathbf{r}) \\ &= (1 + \epsilon\beta)^{-Td} \cdot \frac{1}{4^d} (\epsilon\sigma^{-2})^{(T+1)d} \end{aligned}$$

Thus,

$$\begin{aligned} G(\theta, \theta^*) &= \mathbf{E} \int \mathbf{P} \left( \theta_0, \theta_1, \dots, \theta_{T-1}, \theta^* \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \min(1, a(\boldsymbol{\theta})) d\theta_0 \cdots d\theta_{T-1} \\ &= (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \mathbf{E} \int \mathbf{P} \left( r, \eta_0, \dots, \eta_{T-1} \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \\ &\quad \min(1, a(\boldsymbol{\eta}, r)) d\theta_0 \cdots d\theta_{T-1} \end{aligned}$$

By the distribution of  $r$  and  $\eta_t$ , we know that

$$\begin{aligned} &\mathbf{P} \left( r, \eta_0, \eta_2, \dots, \eta_{T-2} \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \\ &= (2\pi\sigma^2)^{-\frac{d}{2}} \cdot \exp \left( -\frac{\|r\|^2}{2\sigma^2} \right) \cdot \prod_{t=0}^{T-1} (8\pi\epsilon\beta\sigma^2)^{-\frac{d}{2}} \cdot \exp \left( -\frac{\|\eta_t\|^2}{8\epsilon\beta\sigma^2} \right) \\ &= (2\pi\sigma^2)^{-\frac{d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{-\frac{Td}{2}} \cdot \exp \left( -\frac{\|r\|^2}{2\sigma^2} \right) \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \sum_{t=0}^{T-1} \|\eta_t\|^2 \right). \end{aligned}$$

Notice that

$$\begin{aligned}
 \sum_{t=0}^{T-1} \|\eta_t\|^2 &= \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) + \epsilon \beta \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) \right\|^2 \\
 &= \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 \\
 &\quad + 2\epsilon\beta \left( r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right)^T \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2 \\
 &= \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 + 2\epsilon\beta \left( \left\| r_{t+\frac{1}{2}} \right\|^2 - \left\| r_{t-\frac{1}{2}} \right\|^2 \right) \\
 &\quad + 2\epsilon^2 \beta \nabla \tilde{U}_t(\theta_t)^T \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2 \\
 &= 2\epsilon\beta \left( \left\| r_{T-\frac{1}{2}} \right\|^2 - \left\| r \right\|^2 \right) + \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 \\
 &\quad + 4\epsilon\beta\sigma^2 \left( \rho_{t+\frac{1}{2}} - \rho_{t-\frac{1}{2}} \right) + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2 \\
 &= 2\epsilon\beta \left( \left\| r^* \right\|^2 - \left\| r \right\|^2 \right) + 4\epsilon\beta\sigma^2 \left( \rho_{T-\frac{1}{2}} - \rho_{-\frac{1}{2}} \right) \\
 &\quad + \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2.
 \end{aligned}$$

By substituting this above and recalling that  $\rho_{-\frac{1}{2}} = 0$ ,

$$\begin{aligned}
 G(\theta, \theta^*) &= (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \mathbf{E} \int \mathbf{P} \left( r, \eta_0, \eta_2, \dots, \eta_{T-2}, \theta^* \mid \theta, \tilde{U}_0, \dots, \tilde{U}_{T-1} \right) \\
 &\quad \min(1, a(\boldsymbol{\eta}, r)) d\theta_0 \cdots d\theta_{T-1} \\
 &= (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \mathbf{E} \int (2\pi\sigma^2)^{-\frac{d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{-\frac{Td}{2}} \cdot \exp \left( -\frac{\|r\|^2}{2\sigma^2} \right) \\
 &\quad \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \cdot 2\epsilon\beta \left( \left\| r^* \right\|^2 - \left\| r \right\|^2 \right) \right) \\
 &\quad \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \cdot 4\epsilon\beta\sigma^2 \left( \rho_{T-\frac{1}{2}} - \rho_{-\frac{1}{2}} \right) \right) \\
 &\quad \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2 \right) \\
 &\quad \cdot \min(1, a) d\theta_0 \cdots d\theta_{T-1} \\
 &= (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \mathbf{E} \int (2\pi\sigma^2)^{-\frac{d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{-\frac{Td}{2}} \\
 &\quad \cdot \exp \left( -\frac{1}{4\sigma^2} \left( \left\| r^* \right\|^2 + \left\| r \right\|^2 \right) \right) \cdot \exp \left( -\frac{1}{2} \rho_{T-\frac{1}{2}} \right) \\
 &\quad \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left\| r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon \nabla \tilde{U}_t(\theta_t) \right\|^2 + \epsilon^2 \beta^2 \left\| r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right\|^2 \right) \\
 &\quad \cdot \min(1, a) d\theta_0 \cdots d\theta_{T-1}
 \end{aligned}$$

where  $\mathbf{r}$  are to be understood as functions of the  $\theta_t$ , and the integral is taken over  $\theta_t$ .

Substituting the expression of  $a$ , then the term inside the integral is

$$\begin{aligned}
 & (2\pi\sigma^2)^{\frac{-d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{\frac{-Td}{2}} \cdot \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{2}\rho_{T-\frac{1}{2}}\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\right\|^2 + \epsilon^2\beta^2 \left\|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\right\|^2\right) \\
 & \quad \cdot \min\left(1, \exp\left(U(\theta) - U(\theta^*) + \rho_{T-\frac{1}{2}}\right)\right) \\
 & = (2\pi\sigma^2)^{\frac{-d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{\frac{-Td}{2}} \cdot \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\right\|^2 + \epsilon^2\beta^2 \left\|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\right\|^2\right) \\
 & \quad \cdot \exp(U(\theta)) \cdot \min\left(\exp\left(-U(\theta) - \frac{1}{2}\rho_{T-\frac{1}{2}}\right), \exp\left(-U(\theta^*) + \frac{1}{2}\rho_{T-\frac{1}{2}}\right)\right).
 \end{aligned}$$

Finally, this probability multiplied by the probability of  $\theta_0$ , which is  $\frac{1}{Z} \exp(-U(\theta))$ , is

$$\begin{aligned}
 & \pi(\theta)G(\theta, \theta^*) \\
 & = \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \\
 & \quad \mathbf{E} \int (2\pi\sigma^2)^{\frac{-d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{\frac{-Td}{2}} \cdot \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\right\|^2 + \epsilon^2\beta^2 \left\|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\right\|^2\right) \\
 & \quad \cdot \exp(U(\theta)) \cdot \min\left(\exp\left(-U(\theta) - \frac{1}{2}\rho_{T-\frac{1}{2}}\right), \exp\left(-U(\theta^*) + \frac{1}{2}\rho_{T-\frac{1}{2}}\right)\right) d\theta_0 \cdots d\theta_{T-1} \\
 & = \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot 4^d (\epsilon\sigma^{-2})^{-(T+1)d} \cdot (2\pi\sigma^2)^{\frac{-d}{2}} \cdot (8\pi\epsilon\beta\sigma^2)^{\frac{-Td}{2}} \cdot \\
 & \quad \mathbf{E} \int \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left(\left\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\right\|^2 + \epsilon^2\beta^2 \left\|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\right\|^2\right)\right) \\
 & \quad \cdot \min\left(\exp\left(-U(\theta) - \frac{1}{2}\rho_{T-\frac{1}{2}}\right), \exp\left(-U(\theta^*) + \frac{1}{2}\rho_{T-\frac{1}{2}}\right)\right) d\theta_0 \cdots d\theta_{T-1} \\
 & = \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{3(T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \mathbf{E} \int \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left(\left\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\right\|^2 + \epsilon^2\beta^2 \left\|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\right\|^2\right)\right) \\
 & \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \cdot \exp\left(-\frac{1}{2} \left|U(\theta) - U(\theta^*) + \rho_{T-\frac{1}{2}}\right|\right) d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

And writing this out explicitly in terms of

$$\rho_{T-\frac{1}{2}} = \frac{1}{2}\epsilon\sigma^{-2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T (r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}),$$



we get

$$\begin{aligned}
 & \pi(\theta)G(\theta, \theta^*) \\
 &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{3(T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \mathbf{E} \int \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left(\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\|^2 + \epsilon^2\beta^2 \|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\|^2\right)\right) \\
 & \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\
 & \quad \cdot \exp\left(-\frac{1}{2} \left| U(\theta) - U(\theta^*) + \frac{1}{2}\epsilon\sigma^{-2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T (r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}) \right|\right) d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

Now, for this forward path from  $\theta$  to  $\theta^*$ , consider the reverse leapfrog trajectory from  $\theta^*$  to  $\theta$ . This trajectory will have the same values for  $\theta, \theta_0, \dots, \theta^*$  in the reversed order and will have negated values for  $r_{-\frac{1}{2}}, r_{1-\frac{1}{2}}, \dots, r_{T-\frac{1}{2}}$  in the reversed order again. Because of this negation, the values of  $\rho_{\frac{1}{2}}, \rho_{1+\frac{1}{2}}, \dots, \rho_{T-\frac{1}{2}}$  will also be negated. It follows that  $\pi(\theta^*)G(\theta^*, \theta)$  will have the same expression.

Therefore,

$$\pi(\theta)G(\theta, \theta^*) = \pi(\theta^*)G(\theta^*, \theta).$$

This shows that Algorithm 2 with resampling momentum is reversible.

Now we show that the chain satisfies *skew detailed balance* and the stationary distribution of  $\theta$  is  $\pi(\theta)$  if not resampling momentum. Skew detailed balance means that the chain satisfies the following condition (Turitsyn et al., 2011)

$$\pi(x)G(x, y) = \pi(y^\perp)G(y^\perp, x^\perp)$$

where  $G$  is the transition probability.

By Section B, we know that a chain that satisfies the above condition will have invariant distribution  $\pi(x)$ .

In our setting,  $x = (\theta, r)$  and  $x^\perp = (\theta, -r)$ . Given this, the skew detailed balance is

$$\pi(\theta, r)G((\theta, r), (\theta^*, r^*)) = \pi(\theta^*, -r^*)G((\theta^*, -r^*), (\theta, -r)).$$

Next we will show that Algorithm 2 without resampling momentum satisfies the above condition and it naturally follows that Algorithm 2 without resampling converges to the desired distribution.

We consider the joint distribution of  $(\theta, r)$ . By a similar analysis of resampling case, we can get that

$$\begin{aligned}
 & \pi(\theta, r) \cdot G((\theta, r), (\theta^*, r^*)) \\
 &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{3(T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \mathbf{E} \int \exp\left(-\frac{1}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\
 & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left(\|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\|^2 + \epsilon^2\beta^2 \|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\|^2\right)\right) \\
 & \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\
 & \quad \cdot \exp\left(-\frac{1}{2} \left| U(\theta) - U(\theta^*) + \frac{1}{2}\epsilon\sigma^{-2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T (r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}) \right|\right) d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

Again, the reverse trajectory will have the same values for  $\theta, \theta_0, \dots, \theta^*$  and will have negated values for  $r_{-\frac{1}{2}}, r_{1-\frac{1}{2}}, \dots, r_{T-\frac{1}{2}}$  in the reversed order. Therefore,  $\pi(\theta^*, -r^*)G((\theta^*, -r^*), (\theta, -r))$  will have the same expression.

It follows that

$$\pi(\theta, r)G((\theta, r), (\theta^*, r^*)) = \pi(\theta^*, -r^*)G((\theta^*, -r^*), (\theta, -r))$$

which is what we want.  $\square$

## D Connection to HMC

When using a full-batch,  $\beta = 0$ , and resampling, AMAGOLD becomes HMC. To see this, we first notice that with  $\beta = 0$  the update rules of  $\theta$  and  $r$  are the same as in HMC. The remaining thing is to show  $a$  is also the same as in HMC. We rewrite  $\rho_{t+\frac{1}{2}}$  as

$$\rho_{t+\frac{1}{2}} = \rho_{t-\frac{1}{2}} + \frac{1}{2}\sigma^{-2} \left( r_{t-\frac{1}{2}} - r_{t+\frac{1}{2}} \right)^T \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) = \rho_{t-\frac{1}{2}} + \frac{1}{2}\sigma^{-2} \left( \|r_{t-\frac{1}{2}}\|^2 - \|r_{t+\frac{1}{2}}\|^2 \right)$$

As a result,

$$\rho_{T-\frac{1}{2}} = \frac{1}{2}\sigma^{-2} \sum_{t=0}^{T-1} \left( \|r_{t-\frac{1}{2}}\|^2 - \|r_{t+\frac{1}{2}}\|^2 \right) = \frac{1}{2}\sigma^{-2} \left( \|r_{-\frac{1}{2}}\|^2 - \|r_{T-\frac{1}{2}}\|^2 \right)$$

It follows that  $a$  becomes the same as in HMC.

## E Proof of Theorem 2

In this section we prove a bound on the convergence rate of AMAGOLD as compared with second-order Langevin dynamics (L2MC).

*Proof.* We start with the expression we derived for the transition probability in the proof of reversibility.

$$\begin{aligned} & \pi(\theta)G(\theta, \theta^*) \\ &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\ & \quad \cdot \mathbf{E} \int \exp \left( -\frac{1}{4\sigma^2} \left( \|r^*\|^2 + \|r\|^2 \right) \right) \\ & \quad \cdot \exp \left( -\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left( \|r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} + \epsilon\nabla\tilde{U}_t(\theta_t)\|^2 + c^2\epsilon^2\beta^2 \|r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}}\|^2 \right) \right) \\ & \quad \cdot \exp \left( -\frac{U(\theta) + U(\theta^*)}{2} \right) \\ & \quad \cdot \exp \left( -\frac{1}{2} \left| U(\theta) - U(\theta^*) + \frac{1}{2}\epsilon\sigma^{-2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T \left( r_{t-\frac{1}{2}} + r_{t+\frac{1}{2}} \right) \right| \right) d\theta_0 \cdots d\theta_{T-1}. \end{aligned}$$

Since

$$\theta_t = \theta_{t-1} + \epsilon\sigma^{-2}r_{t-\frac{1}{2}},$$

if we define  $\theta_{-1}$  and  $\theta_T$  by convention such that

$$\frac{\theta_{-1} + \theta_0}{2} = \theta \quad \text{and} \quad \frac{\theta_{T-1} + \theta_T}{2} = \theta^*,$$

it follows that for all  $t \in \{0, \dots, T-1\}$

$$r_{t+\frac{1}{2}} - r_{t-\frac{1}{2}} = \epsilon^{-1}\sigma^2 (\theta_{t+1} - 2\theta_t + \theta_{t-1})$$

and

$$r_{t+\frac{1}{2}} + r_{t-\frac{1}{2}} \epsilon^{-1} \sigma^2 (\theta_{t+1} - \theta_{t-1})$$

so we can write the above transition probability explicitly in terms of the  $\theta_t$  as

$$\begin{aligned} & \pi(\theta)G(\theta, \theta^*) \\ &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\ & \quad \cdot \mathbf{E} \int \exp\left(-\frac{c}{4\sigma^2} (\|r^*\|^2 + \|r\|^2)\right) \\ & \quad \cdot \exp\left(-\frac{1}{8\epsilon\beta\sigma^2} \cdot \sum_{t=0}^{T-1} \left(\|\epsilon^{-1}\sigma^2(\theta_{t+1} - 2\theta_t + \theta_{t-1}) + \epsilon\nabla\tilde{U}_t(\theta_t)\|^2\right.\right. \\ & \quad \left.\left.+ c^2\epsilon^2\beta^2\|\epsilon^{-1}\sigma^2(\theta_{t+1} - \theta_{t-1})\|^2\right)\right) \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\ & \quad \cdot \exp\left(-\frac{1}{2} \left|U(\theta) - U(\theta^*) + \frac{1}{2}\epsilon\sigma^{-2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T (\epsilon^{-1}\sigma^2(\theta_{t+1} - \theta_{t-1}))\right|\right) \\ & \quad \cdot d\theta_0 \cdots d\theta_{T-1}. \end{aligned}$$

Simplifying this a bit, we get

$$\begin{aligned} & \pi(\theta)G(\theta, \theta^*) \\ &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\ & \quad \cdot \mathbf{E} \int \exp\left(-\frac{\sigma^2}{\epsilon^2} (\|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2)\right) \\ & \quad \cdot \exp\left(-\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \left\|\theta_{t+1} - 2\theta_t + \theta_{t-1} + \epsilon^2\sigma^{-2}\nabla\tilde{U}_t(\theta_t)\right\|^2\right) \\ & \quad \cdot \exp\left(-\frac{\beta\sigma^2}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|\theta_{t+1} - \theta_{t-1}\|^2\right) \\ & \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\ & \quad \cdot \exp\left(-\frac{1}{2} \left|U(\theta) - U(\theta^*) + \frac{1}{2} \sum_{t=0}^{T-1} \nabla\tilde{U}_t(\theta_t)^T (\theta_{t+1} - \theta_{t-1})\right|\right) d\theta_0 \cdots d\theta_{T-1}. \end{aligned}$$

Next, let

$$\begin{aligned} N_t &= \nabla\tilde{U}_t(\theta_t) - \nabla U_t(\theta_t), \\ A_t &= \theta_{t+1} - 2\theta_t + \theta_{t-1} + \epsilon^2\sigma^{-2}\nabla U_t(\theta_t), \\ B_t &= \theta_{t+1} - \theta_{t-1}, \\ C_t &= U(\theta) - U(\theta^*) + \frac{1}{2} \sum_{t=0}^{T-1} \nabla U_t(\theta_t)^T (\theta_{t+1} - \theta_{t-1}). \end{aligned}$$

Notice that only  $N_t$  depends on the randomness of the stochastic gradient samples. Then,

$$\begin{aligned}
 & \pi(\theta)G(\theta, \theta^*) \\
 &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \int \mathbf{E} \left[ \exp \left( -\frac{\sigma^2}{\epsilon^2} \left( \|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2 \right) \right) \right. \\
 & \quad \cdot \exp \left( -\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \left( \|A_t\|^2 + 2\epsilon^2\sigma^{-2}A_t^T N_t + \epsilon^4\sigma^{-4}\|N_t\|^2 \right) \right) \\
 & \quad \cdot \exp \left( -\frac{\beta\sigma^2}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|B_t\|^2 \right) \\
 & \quad \cdot \exp \left( -\frac{U(\theta) + U(\theta^*)}{2} \right) \\
 & \quad \left. \cdot \exp \left( -\frac{1}{2} \left| C_t + \frac{1}{2} \sum_{t=0}^{T-1} N_t^T B_t \right| \right) \right] d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

Now, for any constant  $c > 1$ , we can bound

$$\begin{aligned}
 \mathbf{E} \left[ \left\| \sum_{t=0}^{T-1} N_t^T B_t \right\| \right] &\leq \sqrt{\mathbf{E} \left[ \left( \sum_{t=0}^{T-1} N_t^T B_t \right)^2 \right]} \\
 &= \sqrt{\sum_{t=0}^{T-1} B_t^T \mathbf{E} [N_t N_t^T] B_t} \\
 &\leq \sqrt{\sum_{t=0}^{T-1} \frac{V^2}{d} \|B_t\|^2} \\
 &\leq \frac{V^2}{d} \frac{\epsilon}{2(c-1)\beta\sigma^2} + (c-1) \frac{\beta\sigma^2}{2\epsilon} \sum_{t=0}^{T-1} \|B_t\|^2.
 \end{aligned}$$

Additionally, we know that  $\mathbf{E} [N_t] = 0$  and

$$\mathbf{E} [\|N_t\|^2] = \mathbf{E} [\text{tr} (N_t N_t^T)] \leq \text{tr} \left( \frac{V^2}{d} I \right) = V^2.$$

So, since by Jensen's inequality,  $\mathbf{E} [\exp(X)] \geq \exp(\mathbf{E} [X])$ , we can bound this with

$$\begin{aligned}
 & \pi(\theta)G(\theta, \theta^*) \\
 &\geq \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \int \exp \left( -\frac{\sigma^2}{\epsilon^2} \left( \|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2 \right) \right) \\
 & \quad \cdot \exp \left( -\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \|A_t\|^2 \right) \cdot \exp \left( -\frac{\epsilon TV^2}{8\sigma^2\beta} \right) \\
 & \quad \cdot \exp \left( -\frac{c\sigma^2\beta}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|B_t\|^2 \right) \\
 & \quad \cdot \exp \left( -\frac{U(\theta) + U(\theta^*)}{2} \right) \\
 & \quad \cdot \exp \left( -\frac{1}{2} |C_t| \right) \cdot \exp \left( -\frac{1}{(c-1)Td} \cdot \frac{\epsilon TV^2}{8\sigma^2\beta} \right) d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

Now, this is a lower bound on the AMAGOLD chain with parameters  $(\epsilon, \sigma, \beta)$ . Next, we consider the transition probability of a *rescaled* chain, with slightly different parameters, that will be set as a function of  $c$ . Specifically, consider the chain with parameters  $(\epsilon, \sigma \cdot c^{-1/4}, \beta \cdot c^{-1/2})$ . (We will set the parameter  $c$  later; at this point in the proof it is just an arbitrary constant  $c > 1$ .) If we call this rescaled chain  $G_r$ , then by substitution of the parameters into the above expression, we get

$$\begin{aligned}
& \pi(\theta)G_r(\theta, \theta^*) \\
& \geq \frac{1}{Z} \cdot (1 + c^{-1/2}\epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{3(T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot c^{-\frac{d}{4}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
& \quad \cdot \int \mathbf{E} \exp\left(-\frac{\sigma^2}{c^{1/2}\epsilon^2} \left(\|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2\right)\right) \\
& \quad \cdot \exp\left(-\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \|A_t\|^2\right) \cdot \exp\left(-\frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \\
& \quad \cdot \exp\left(-\frac{\sigma^2\beta}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|B_t\|^2\right) \\
& \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\
& \quad \cdot \exp\left(-\frac{1}{2}|C_t|\right) \cdot \exp\left(-\frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right) d\theta_0 \cdots d\theta_{T-1} \\
& \geq \frac{1}{Z} \cdot (1 + c^{-1/2}\epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{3(T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot c^{-\frac{d}{4}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
& \quad \cdot \exp\left(-\frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \cdot \exp\left(-\frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \\
& \quad \cdot \int \mathbf{E} \exp\left(-\frac{\sigma^2}{\epsilon^2} \left(\|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2\right)\right) \\
& \quad \cdot \exp\left(-\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \|A_t\|^2\right) \\
& \quad \cdot \exp\left(-\frac{\sigma^2\beta}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|B_t\|^2\right) \\
& \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\
& \quad \cdot \exp\left(-\frac{1}{2}|C_t|\right) d\theta_0 \cdots d\theta_{T-1}.
\end{aligned}$$

On the other hand, consider the transition probability of the full-gradient L2MC chain with parameters  $(\epsilon, \sigma, \beta)$ . This chain will be the same as the AMAGOLD chain, except that  $N_t = 0$  always. So, if we call this chain's

transition probability  $\bar{G}$ , we will have

$$\begin{aligned}
 & \pi(\theta)\bar{G}(\theta, \theta^*) \\
 &= \frac{1}{Z} \cdot (1 + \epsilon\beta)^{Td} \cdot \beta^{-\frac{Td}{2}} \cdot 2^{-\frac{(3T-1)d}{2}} \cdot \pi^{-\frac{(T+1)d}{2}} \cdot \epsilon^{-\frac{(3T+2)d}{2}} \cdot \sigma^{(T+1)d} \\
 & \quad \cdot \int \exp\left(-\frac{\sigma^2}{\epsilon^2} \left(\|\theta^* - \theta_{T-1}\|^2 + \|\theta_0 - \theta\|^2\right)\right) \\
 & \quad \cdot \exp\left(-\frac{\sigma^2}{8\epsilon^3\beta} \cdot \sum_{t=0}^{T-1} \|A_t\|^2\right) \\
 & \quad \cdot \exp\left(-\frac{\beta\sigma^2}{8\epsilon} \cdot \sum_{t=0}^{T-1} \|B_t\|^2\right) \\
 & \quad \cdot \exp\left(-\frac{U(\theta) + U(\theta^*)}{2}\right) \\
 & \quad \cdot \exp\left(-\frac{1}{2}|C_t|\right) d\theta_0 \cdots d\theta_{T-1}.
 \end{aligned}$$

Using this, we can simplify our bound on the transition probability of the AMAGOLD chain to

$$\begin{aligned}
 \pi(\theta)G_r(\theta, \theta^*) &\geq \left(\frac{1 + c^{-1/2}\epsilon\beta}{1 + \epsilon\beta}\right)^{Td} \cdot c^{-\frac{d}{4}} \cdot \exp\left(-\frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \cdot \exp\left(-\frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \\
 &\quad \cdot \pi(\theta)\bar{G}(\theta, \theta^*).
 \end{aligned}$$

Thus,

$$\frac{\pi(\theta)G_r(\theta, \theta^*)}{\pi(\theta)\bar{G}(\theta, \theta^*)} \geq \left(\frac{1 + c^{-1/2}\epsilon\beta}{1 + \epsilon\beta}\right)^{Td} \cdot c^{-\frac{d}{4}} \cdot \exp\left(-\frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \cdot \exp\left(-\frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right).$$

All that remains to get a bound is to set  $c$  appropriately. Since

$$c^{-\frac{d}{4}} = \exp\left(-\frac{d}{4}\log(c)\right) \geq \exp\left(-\frac{d}{4}(c-1)\right),$$

and

$$c^{-1/2} \geq 1 - \frac{c-1}{2},$$

we can bound this with

$$\begin{aligned}
 \frac{\pi(\theta)G_r(\theta, \theta^*)}{\pi(\theta)\bar{G}(\theta, \theta^*)} &\geq \left(\frac{1 + (1 - \frac{c-1}{2})\epsilon\beta}{1 + \epsilon\beta}\right)^{Td} \cdot \exp\left(-\frac{d}{4}(c-1) - \frac{c\epsilon TV^2}{8\sigma^2\beta} - \frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \\
 &\geq \left(1 - \frac{(c-1)\epsilon\beta}{2(1 + \epsilon\beta)}\right)^{Td} \cdot \exp\left(-\frac{d}{4}(c-1) - \frac{c\epsilon TV^2}{8\sigma^2\beta} - \frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right).
 \end{aligned}$$

Since for any  $0 \leq x < 1/2$ , it holds that  $1 - x \geq \exp(-2x)$ , as long as

$$\frac{(c-1)\epsilon\beta}{1 + \epsilon\beta} \leq 1,$$

it holds that

$$1 - \frac{(c-1)\epsilon\beta}{2(1 + \epsilon\beta)} \geq \exp\left(-\frac{(c-1)\epsilon\beta}{1 + \epsilon\beta}\right).$$

So, under this assumption,

$$\begin{aligned}
 \frac{\pi(\theta)G_r(\theta, \theta^*)}{\pi(\theta)\bar{G}(\theta, \theta^*)} &\geq \exp\left(-\frac{(c-1)\epsilon\beta Td}{1+\epsilon\beta} - \frac{d}{4}(c-1) - \frac{c\epsilon TV^2}{8\sigma^2\beta} - \frac{1}{(c-1)Td} \cdot \frac{c\epsilon TV^2}{8\sigma^2\beta}\right) \\
 &= \exp\left(-\frac{\epsilon TV^2}{8\sigma^2\beta} - \frac{\epsilon V^2}{8\sigma^2\beta d}\right) \\
 &\quad \cdot \exp\left(-(c-1)\left(\frac{(1+\epsilon\beta(1+4T))d}{4(1+\epsilon\beta)} + \frac{\epsilon TV^2}{8\sigma^2\beta}\right) - \frac{\epsilon V^2}{8(c-1)\sigma^2\beta d}\right) \\
 &= \exp\left(-\frac{\epsilon TV^2}{8\sigma^2\beta} - \frac{\epsilon V^2}{8\sigma^2\beta d}\right) \\
 &\quad \cdot \exp\left(-(c-1)\left(\frac{3}{2}Td + \frac{\epsilon TV^2}{8\sigma^2\beta}\right) - \frac{\epsilon V^2}{8(c-1)\sigma^2\beta d}\right).
 \end{aligned}$$

If we also assume that

$$\frac{\epsilon V^2}{4\sigma^2\beta d} \leq 1,$$

then

$$\frac{\epsilon TV^2}{8\sigma^2\beta} \leq \frac{Td}{2},$$

and so

$$\begin{aligned}
 \frac{\pi(\theta)G_r(\theta, \theta^*)}{\pi(\theta)\bar{G}(\theta, \theta^*)} &\geq \exp\left(-\frac{\epsilon TV^2}{8\sigma^2\beta} - \frac{\epsilon V^2}{8\sigma^2\beta d}\right) \\
 &\quad \cdot \exp\left(-(c-1)2Td - \frac{\epsilon V^2}{8(c-1)\sigma^2\beta d}\right).
 \end{aligned}$$

Next, set

$$c-1 = \sqrt{\frac{\epsilon V^2}{16\sigma^2\beta Td^2}}.$$

From this, we will get

$$\begin{aligned}
 \frac{\pi(\theta)G_r(\theta, \theta^*)}{\pi(\theta)\bar{G}(\theta, \theta^*)} &\geq \exp\left(-\frac{\epsilon TV^2}{8\sigma^2\beta} - \frac{\epsilon V^2}{8\sigma^2\beta d}\right) \cdot \exp\left(-\sqrt{\frac{\epsilon TV^2}{\sigma^2\beta}}\right) \\
 &\geq \exp\left(-\frac{\epsilon TV^2}{4\sigma^2\beta} - \sqrt{\frac{\epsilon TV^2}{\sigma^2\beta}}\right).
 \end{aligned}$$

Now, in order for this to hold, we needed

$$\frac{(c-1)\epsilon\beta}{1+\epsilon\beta} \leq 1.$$

With our setting of  $c$ , and our other assumption,

$$c-1 = \sqrt{\frac{\epsilon V^2}{16\sigma^2\beta Td^2}} = \sqrt{\frac{\epsilon V^2}{4\sigma^2\beta d} \cdot \frac{1}{4Td}} \leq \sqrt{\frac{1}{4Td}} \leq 1,$$

so the bound will trivially hold. Thus the only added assumption we needed is the one stated in the Theorem statement, that

$$\frac{\epsilon V^2}{4\sigma^2\beta d} \leq 1.$$

Now we apply the standard Dirichlet form argument. The spectral gap of a Markov chain can be written as (Aida, 1998)

$$\gamma = \inf_{f \in L_0^2(\pi): \text{Var}_\pi[f]=1} \mathcal{E}(f)$$

where  $L_0^2(\pi)$  denotes the Hilbert space of all functions that are square integrable with respect to probability measure  $\pi$  and have mean zero.  $\mathcal{E}(f)$  is the Dirichlet form of a Markov chain associated with transition operator  $T$  (Fukushima et al., 2010):

$$\mathcal{E}(f) = \frac{1}{2} \int \int [(f(\theta) - f(\theta^*))^2] G(\theta, \theta^*) \pi(\theta) d\theta d\theta^*$$

By the expression of the spectral gap, it follows that

$$\begin{aligned} \gamma &= \inf_{f \in L_0^2(\pi): \text{Var}_\pi[f]=1} \left[ \frac{1}{2} \int \int [(f(\theta) - f(\theta^*))^2] G(\theta, \theta^*) \pi(\theta) d\theta d\theta^* \right] \\ &\geq \exp\left(-\frac{\epsilon TV^2}{4\sigma^2\beta} - \sqrt{\frac{\epsilon TV^2}{\sigma^2\beta}}\right) \cdot \inf_{f \in L_0^2(\pi): \text{Var}_\pi[f]=1} \left[ \frac{1}{2} \int \int [(f(\theta) - f(\theta^*))^2] \bar{G}(\theta, \theta) \pi(\theta) d\theta d\theta^* \right] \\ &= \exp\left(-\frac{\epsilon TV^2}{4\sigma^2\beta} - \sqrt{\frac{\epsilon TV^2}{\sigma^2\beta}}\right) \cdot \bar{\gamma} \end{aligned}$$

This finishes the proof. □

## F Reformulation of AMAGOLD Algorithm

We reformulate our algorithm by setting  $v = \epsilon\sigma^{-2}r, b = \epsilon\beta, h = \epsilon^2\sigma^{-2}$  and outline the algorithm after reformulation in Algorithm 3.

---

### Algorithm 3 Reformulated AMAGOLD

---

- 1: **given:** Energy  $U$ , initial state  $\theta \in \Theta$
  - 2: **loop**
  - 3:   **optionally, resample momentum:**  $v \sim \mathcal{N}(0, h\mathbf{I})$
  - 4:   **initialize momentum and energy acc:**  $v_{-\frac{1}{2}} \leftarrow v, \rho_{-\frac{1}{2}} \leftarrow 0$
  - 5:   **half position update:**  $\theta_0 \leftarrow \theta + \frac{1}{2}v_{-\frac{1}{2}}$
  - 6:   **for**  $t = 0$  **to**  $T - 1$  **do**
  - 7:     **if**  $t \neq 0$  **then**
  - 8:       **position update:**  $\theta_t \leftarrow \theta_{t-1} + v_{t-\frac{1}{2}}$
  - 9:     **end if**
  - 10:    **sample noise**  $\eta_t \sim \mathcal{N}(0, 4hb)$
  - 11:    **sample random energy component**  $\tilde{U}_t$
  - 12:    **update momentum:**  $v_{t+\frac{1}{2}} \leftarrow \left( (1-b)v_{t-\frac{1}{2}} - h\nabla\tilde{U}_t(\theta_t) + \eta_t \right) / (1+b)$
  - 13:    **update energy acc:**  $\rho_{t+\frac{1}{2}} \leftarrow \rho_{t-\frac{1}{2}} + \frac{1}{2}\nabla\tilde{U}_t(\theta_t)^T \left( v_{t-\frac{1}{2}} + v_{t+\frac{1}{2}} \right)$
  - 14:    **end for**
  - 15:    **half position update:**  $\theta_T \leftarrow \theta_{T-1} + \frac{1}{2}v_{T-\frac{1}{2}}$
  - 16:    **new values:**  $\theta^* \leftarrow \theta_T, v^* \leftarrow v_{T-\frac{1}{2}}$
  - 17:     $a \leftarrow \exp\left(U(\theta) - U(\theta^*) + \rho_{T-\frac{1}{2}}\right)$
  - 18:    **with probability**  $\min(1, a)$  **update**  $\theta \leftarrow \theta^*, v \leftarrow v^*$  (as long as  $\theta^* \in \Theta$ )
  - 19:    **otherwise update**  $v \leftarrow -v_{-\frac{1}{2}}$
  - 20: **end loop**
-



## G Additional Experiments Results and Setting Details

### G.1 Double Well Potential

We visualize the estimated density on additional step size settings. Consistent with Figure 1d, it is clear here that SGHMC is very sensitive to step size. A small change in step size will cause a big difference in the estimated density. In contrast, AMAGOLD is more robust and can work well with a large range of step sizes.

When the setup of step size is inappropriate, as in Figures 6a and b where it is fixed to be too small, either SGHMC or AMAGOLD converges in the training time. This is because the chain moves too slowly toward the stationary distribution. However, AMAGOLD with step size tuning is able to automatically adjust the step size based on the information provided by M-H step. As shown in Figure 1c, tuned AMAGOLD can determine a step size that causes convergence given the same training time budget. All results are obtained by collecting  $10^5$  samples with 1000 burn-in samples.

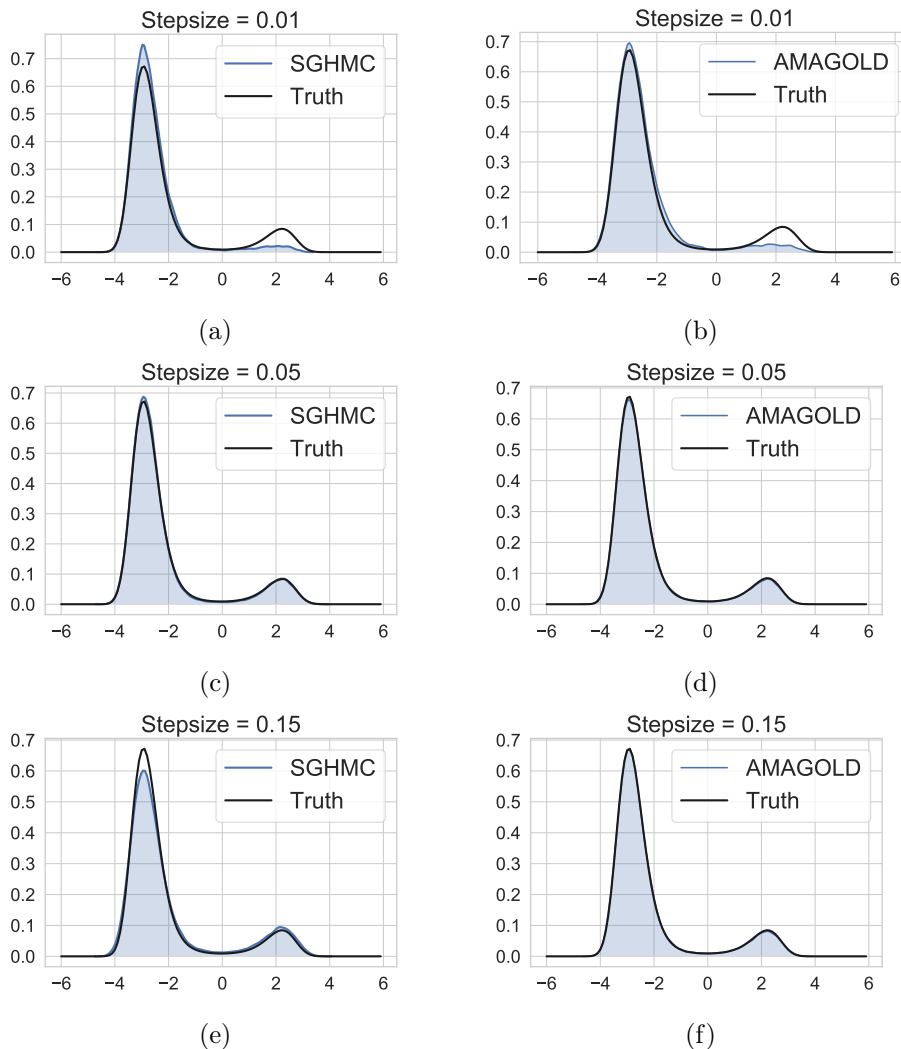


Figure 6: Estimated densities of SGHMC (1st column) and AMAGOLD (2nd column) on varying step sizes.

## G.2 Two-Dimensional Synthetic Distributions

### G.2.1 Analytical Expression

$$\text{Dist1: } \mathcal{N}(z_1; z_2^2/4, 1)\mathcal{N}(z_2; 0, 4)$$

$$\text{Dist2: } 0.5\mathcal{N}\left(z; 0, \begin{bmatrix} 2 & 1.8 \\ 1.8 & 2 \end{bmatrix}\right) + 0.5\mathcal{N}\left(z; 0, \begin{bmatrix} 2 & -1.8 \\ -1.8 & 2 \end{bmatrix}\right)$$

### G.2.2 Runtime Comparisons

We report runtime comparisons between AMAGOLD and SGHMC on Dist1 and Dist2 with step size 0.15 (Figure 7). This experiment uses the analytical energy expression (no data examples), so there is no speed-up of stochastic methods over full-batch methods. At the beginning, SGHMC converges faster due to the lack of M-H step, but eventually it converges to a biased distribution. AMAGOLD is not much slower than SGHMC, which shows that AMA can reduce the amount of computation of adding M-H step while keep the chain unbiased.




Figure 7: Runtime comparisons between SGHMC and AMAGOLD on synthetic distributions (a) Dist1 and (b) Dist2.

### G.2.3 Additional Note on Figure 2

It is worth noting that, even though it is lower than SGHMC’s, AMAGOLD’s KL divergence grows when the step size is large compared to full-batch methods. This is because the M-H acceptance probability decreases, causing the chain to converge more slowly. This is expected. It is well-known that stochastic methods are more sensitive to step sizes than full-batch methods (Nemirovski et al., 2009). However, since AMAGOLD’s KL divergence grows much slower than SGHMC’s, AMAGOLD is more robust to different step sizes

## G.3 Bayesian Logistic Regression

We report the acceptance probability of AMAGOLD on *Heart* for varying step sizes in Figure 8. For a large range of step sizes, the acceptance rate is sufficiently high to allow the chain converge fast, demonstrated in Figure 4. The acceptance rate may become very low with a large step size resulting in slow move. But this undesired acceptance probability can be easily detected and avoided in practice.

## G.4 Bayesian Neural Networks

The architecture of Bayesian Neural Networks is a two-layer MLP with first hidden layer size 500 and the second hidden layer size 256.




Figure 8: The acceptance probability of the M-H step in AMAGOLD for varying of step sizes on the Heart dataset.