
Adaptive Tuning for Metropolis Adjusted Langevin Trajectories

Lionel Riou-Durand
University of Warwick

Pavel Sountsov
Google Research

Jure Vogrinc
University of Warwick

Charles C. Margossian
Flatiron Institute

Sam Power
University of Bristol

Abstract

Hamiltonian Monte Carlo (HMC) is a widely used sampler for continuous probability distributions. In many cases, the underlying Hamiltonian dynamics exhibit a phenomenon of resonance which decreases the efficiency of the algorithm and makes it very sensitive to hyperparameter values. This issue can be tackled efficiently, either via the use of trajectory length randomization (RHMC) or via partial momentum refreshment. The second approach is connected to the kinetic Langevin diffusion, and has been mostly investigated through the use of Generalized HMC (GHMC). However, GHMC induces momentum flips upon rejections causing the sampler to backtrack and waste computational resources. In this work we focus on a recent algorithm bypassing this issue, named Metropolis Adjusted Langevin Trajectories (MALT). We build upon recent strategies for tuning the hyperparameters of RHMC which target a bound on the Effective Sample Size (ESS) and adapt it to MALT, thereby enabling the first user-friendly deployment of this algorithm. We construct a method to optimize a sharper bound on the ESS and reduce the estimator variance. Easily compatible with parallel implementation, the resultant Adaptive MALT algorithm is competitive in terms of ESS rate and hits useful tradeoffs in memory usage when compared to GHMC, RHMC and NUTS.

1 INTRODUCTION

Hamiltonian Monte Carlo (HMC; [Duane et al. \(1987\)](#); [Neal \(2011\)](#)) is a widely-known MCMC sampler for distributions with differentiable densities. The use of multiple gradient-informed steps per proposal suppresses random walk behavior and makes HMC particularly powerful for high-dimensional numerical integration. Its implementation in probabilistic-programming languages (e.g. [Carpenter et al., 2017](#); [Bingham et al., 2018](#); [Salvatier et al., 2016](#); [Lao et al., 2020](#)) leverages automatic differentiation and adaptive tuning with the no-U-turn sampler (NUTS; [Hoffman and Gelman \(2014\)](#)) which require minimal input from the practitioner.

Robust implementations of HMC rely on randomizing the length of Hamiltonian trajectories ([Bou-Rabee and Sanz-Serna, 2017](#), RHMC) to avoid resonances. NUTS is well established and remains wildly popular but recent studies suggest optimally tuned RHMC samplers can perform better. Furthermore, NUTS’ recursive construction (control-flow-heavy) complicates its implementation when running many parallel chains ([Radul et al., 2020](#); [Lao and Dillon, 2019](#); [Phan and Pradhan, 2019](#)). To take better advantage of modern computing hardware such as GPUs, several schemes for tuning RHMC have been proposed ([Hoffman et al., 2021](#); [Sountsov and Hoffman, 2021](#)), showing significant gains of efficiency compared to NUTS.

A second approach to mitigate resonances is based on the kinetic Langevin diffusion. Built upon its discretization, Generalized HMC ([Horowitz, 1991](#), GHMC) alternates single adjusted steps with partial momentum refreshments. However, the original GHMC algorithm had not been proven competitive, its efficiency being restrained by momentum flips induced upon rejections. To reduce the backtracking caused by these flips, [Neal \(2020\)](#) proposed a slice-sampling approach to cluster the event times of the rejections of a proposal. The efficiency of GHMC was significantly improved by this method, which was later turned into a tuning-free algorithm ([Hoffman and Sountsov, 2022](#)).

In this paper, we focus on a new numerical scheme for approximating the kinetic Langevin diffusion, recently introduced as Metropolis Adjusted Langevin Trajectories (Riou-Durand and Vogrinc, 2022, MALT). By construction, MALT enables full erasing of the momentum flips and shares several desirable properties with RHMC. Compared to GHMC, some qualitative arguments in favour of MALT were made but no quantitative comparison had been conducted. One objective is to fill this gap.

Our general goal is to develop a tuning-free version of the MALT algorithm, compatible with running many parallel chains. In Section 3 we develop tuning heuristics and algorithmic solutions for each parameter. Built upon these principles, we present in Section 4 an adaptive sampler, applicable for multiple chains as well as for single chain settings. In Section 5 we compare numerically its efficiency to several adaptive samplers (GHMC, RHMC, NUTS) on a set of benchmark models.

To summarize our contributions:

- We develop tuning heuristics for each of MALT’s hyperparameters, building upon recent tuning strategies for RHMC (Hoffman et al., 2021; Sountsov and Hoffman, 2021) which target a bound on the Effective Sample Size (ESS).
- We propose a method to optimize a sharper bound on the ESS, for which we derive a gradient estimate with reduced variance. This variance reduction method is applicable to both MALT and RHMC and allows for speeding up the adaptation of the trajectory length.
- We construct an Adaptive MALT algorithm, easily compatible with, but not reliant on, running multiple chains in parallel.
- We show that Adaptive MALT is competitive in terms of ESS rate and hits useful tradeoffs in memory usage when compared to previous adaptive samplers (GHMC, RHMC, NUTS).

2 BACKGROUND

This section reviews the use of Hamiltonian dynamics to construct MCMC samplers targeting a differentiable density $\Pi(x) \propto \exp(-\Phi(x))$, for $x \in \mathbb{R}^d$. We introduce the notion of partial momentum refreshment and present its connection to the kinetic Langevin diffusion, thereby motivating GHMC. Unfortunately the utility of GHMC is hindered by momentum flips, required to preserve detailed balance but statistically inefficient; we explain how MALT bypasses this issue. Finally we review existing tuning strategies for HMC and GHMC samplers.

2.1 Hamiltonian Dynamics

To describe the motion of a particle $x \in \mathbb{R}^d$ and its velocity $v \in \mathbb{R}^d$ (a.k.a. momentum), Hamiltonian dynamics are defined for a positive definite mass matrix $M \in \mathbb{R}^{d \times d}$ as

$$dx_t = M^{-1}v_t dt, \quad dv_t = -\nabla\Phi(x_t) dt. \quad (1)$$

By construction, Hamiltonian dynamics preserve the density $\Pi_*(x, v) \propto \exp(-\Phi(x) - v^\top M^{-1}v/2)$. In general, the ODE (1) does not have a closed form solution. Instead, the leapfrog method is commonly used to integrate this ODE numerically. It is defined for a time-step $h > 0$ as the following update $(x_1, v_1) = \text{LEAPFROG}(x_0, v_0; M, h)$ such that

$$\begin{aligned} v_{1/2} &= v_0 - (h/2)\nabla\Phi(x_0) \\ x_1 &= x_0 + hM^{-1}v_{1/2} \\ v_1 &= v_{1/2} - (h/2)\nabla\Phi(x_1). \end{aligned} \quad (2)$$

The leapfrog integrator is time-reversible (Neal, 2011), i.e. $(x_0, -v_0) = \text{LEAPFROG}(x_1, -v_1; M, h)$. This means that going backwards in time is equivalent to flipping the momenta. HMC is then built upon the following recursion. From a position $x_0 \in \mathbb{R}^d$ and a velocity $v_0 \sim \mathcal{N}_d(0_d, M)$, a Hamiltonian trajectory of length $\tau > 0$ is proposed by composing the leapfrog update for $L = \lceil \tau/h \rceil$ steps. The resultant proposal (x_L, v_L) is faced with an accept-reject test (a.k.a. Metropolis correction) to compensate for the numerical error when solving (1). We note $|v|_{M^{-1}}^2 = v^\top M^{-1}v$. With probability $1 - \exp(-\Delta^+)$, where $\Delta^+ = \max(0, \Delta)$ and

$$\Delta = \Phi(x_L) - \Phi(x_0) + (|v_L|_{M^{-1}}^2 - |v_0|_{M^{-1}}^2)/2,$$

the proposal is rejected and its momentum is flipped i.e. the output is $(x_0, -v_0)$. Flipping the momentum upon rejection is a technical requirement to ensure that the Metropolis correction preserves Π_* . Overall, the momentum flips are erased since v_0 is fully refreshed at each iteration.

HMC prevents random walk behaviour by proposing trajectories with persistent momentum, thereby exploring the sampling space more efficiently. However, HMC exhibits a phenomenon of resonance when the target Π has heterogeneous scales. Typically, fast mixing for one component can result in arbitrarily slow mixing for other components (Riou-Durand and Vogrinc, 2022). This issue makes the efficiency of HMC highly sensitive to the choice of its tuning parameters.

One solution consists on drawing τ at random at each iteration (RHMC), thereby smoothing the auto-correlations between successive Hamiltonian trajectories (Bou-Rabee and Sanz-Serna, 2017). The choice of the sampling distribution for τ is further discussed in Section 5. In this work, we consider another approach based on partial momentum refreshment.

2.2 Partial Momentum Refreshment

The resonances of Hamiltonian trajectories can be reduced by refreshing partially the velocity along their path. Kinetic Langevin dynamics leverage this property by augmenting (1) with a continuous momentum refreshment induced by a Brownian motion $(B_t)_{t \geq 0}$ on \mathbb{R}^d . These are defined for a damping parameter $\gamma \geq 0$ (a.k.a. friction), as

$$\begin{aligned} dx_t &= M^{-1}v_t dt, \\ dv_t &= -(\nabla\Phi(x_t) + \gamma v_t) dt + \sqrt{2\gamma}M^{1/2}dB_t. \end{aligned} \quad (3)$$

The damping controls the refreshment speed. As $\gamma \rightarrow 0$, the SDE (3) reduces to the ODE (1) and exhibits resonances. These resonances vanish as γ gets larger. Increasing γ too much however decreases the persistence of the momentum and encourages random walk behavior. This tradeoff can be solved for a critical choice of damping. For log-concave distributions, quantitative mixing rates have been established by Eberle et al. (2019); Dalalyan and Riou-Durand (2020); Cao et al. (2019). Choosing a mass M in (3) is equivalent to preconditioning the augmented space with the linear map $(x, v) \mapsto (M^{1/2}x, M^{-1/2}v)$.

Langevin trajectories can be approximated numerically by alternating between the leapfrog update (2) and a partial momentum refreshment, defined for $\eta = e^{-\gamma h}$ as

$$v' = \eta v + \sqrt{1 - \eta^2} \xi, \quad \xi \sim \mathcal{N}_d(0_d, M).$$

This partial momentum refreshment leaves Π_* invariant. To compensate for the numerical error, one strategy consists in applying the Metropolis correction to each leapfrog step. GHMC is built upon this principle (Horowitz, 1991), thereby inducing a momentum flip to every rejected leapfrog step; see Section 2.1. Unlike HMC, momentum flips are not erased with GHMC, which causes the sampler to backtrack and limits its efficiency (Kennedy and Pendleton, 2001). This also exacerbates the problem of tuning h , which must be small enough to prevent momentum flips from becoming overwhelming (Bou-Rabee and Vanden-Eijnden, 2010).

Several solutions have been proposed to improve GHMC's efficiency. Delayed rejection methods have been used to reduce the number of momentum flips at a higher computational cost (Mira, 2001; Campos and Sanz-Serna, 2015; Park and Atchadé, 2020). We highlight a recent approach based on slice sampling (Neal, 2020), which correlates the accept-reject coin tosses over time and encourages larger periods without momentum flips.

2.3 Metropolis Adjusted Langevin Trajectories

In this work, we focus on a sampler recently proposed by Riou-Durand and Vogrinc (2022) to bypass the problem caused by momentum flips. This method consists in

applying the Metropolis correction to an entire numerical Langevin trajectory, rather than applying it to each of its $L = \lceil \tau/h \rceil$ leapfrog steps. The resultant sampler, named Metropolis Adjusted Langevin Trajectories (MALT), is defined by iterating Algorithm 1.

Algorithm 1 Metropolis Adjusted Langevin Trajectory

```

1: function MALT( $X; M, \gamma, h, \tau$ )
2:   Set  $L \leftarrow \lceil \tau/h \rceil$  and  $\eta \leftarrow e^{-\gamma h}$ 
3:   Draw independent  $(\xi_i)_{i=0, \dots, L} \sim \mathcal{N}_d(0_d, M)$ 
4:   Set  $(x_0, v_0) \leftarrow (X, \xi_0)$  and  $\Delta \leftarrow 0$ 
5:   for  $i = 1$  to  $L$  do
6:     Refresh  $v'_{i-1} \leftarrow \eta v_{i-1} + \sqrt{1 - \eta^2} \xi_i$ 
7:     Set  $(x_i, v_i) \leftarrow \text{LEAPFROG}(x_{i-1}, v'_{i-1}; M, h)$ 
8:     Update  $\Delta \leftarrow \Delta + (|v_i|_{M^{-1}}^2 - |v'_{i-1}|_{M^{-1}}^2)/2$ 
9:   end for
10:  Set  $X' \leftarrow x_L$  and  $\Delta \leftarrow \Delta + \Phi(x_L) - \Phi(x_0)$ 
11:  Draw  $Z \sim \mathcal{Exp}(1)$ 
12:  if  $Z < \Delta$  then set  $X' \leftarrow X$ 
13:  end if
14:  return  $X', v_L, x_0, v'_0, \Delta$ 
15: end function

```

In this algorithm, the overall numerical error unfolds as $\Delta = \sum_{i=1}^L \delta_i$, where

$$\delta_i = \Phi(x_i) - \Phi(x_{i-1}) + (|v_i|_{M^{-1}}^2 - |v'_{i-1}|_{M^{-1}}^2)/2$$

is the energy difference incurred by the i^{th} leapfrog step. As a consequence, the probability of accepting an entire trajectory for MALT is higher than the probability of accepting L consecutive leapfrog steps for GHMC. Indeed:

$$\mathbb{P}_{\text{GHMC}}^{\text{accept}} = e^{-\sum_{i=1}^L \delta_i^+} \leq e^{-(\sum_{i=1}^L \delta_i)^+} = \mathbb{P}_{\text{MALT}}^{\text{accept}}.$$

Detailed balance and ergodicity of Algorithm 1 are proven in Riou-Durand and Vogrinc (2022). Notably, the momentum flips get erased by fully refreshing $v_0 \sim \mathcal{N}(0_d, M)$ at the start of each trajectory. This ensures reversibility of the sampler with respect to Π , and makes HMC a special case of MALT obtained for $\gamma = 0$. As γ gets larger, Langevin trajectories become ergodic, thereby preventing U-turns and resonances. Compared to GHMC, Algorithm 1 also exhibits an advantage for tuning the step-size, which is supported by optimal scaling results derived in Riou-Durand and Vogrinc (2022). The tuning of MALT's parameters is further discussed in Section 3.

2.4 Adaptive Tuning Strategies

To measure the efficiency of a MCMC sampler, we consider the problem of estimating $\mathbb{E}[\varphi(x)]$ as $x \sim \Pi$ for a square integrable function $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$. The Monte Carlo estimator, $\hat{\varphi}_N = N^{-1} \sum_{n=1}^N \varphi(X_n)$, estimates the desired expectation value, with N the number of approximate samples. A scale-free measurement of $\hat{\varphi}_N$'s precision is given

by the Effective Sample Size (ESS), defined as

$$\text{ESS}_\varphi = \frac{N}{C_\varphi}, \quad C_\varphi = 1 + 2 \sum_{n=1}^{\infty} \text{Corr}(\varphi(X_n), \varphi(X_0))$$

where X_n is the n^{th} iterate of the Markov chain at stationarity. Asymptotically, ESS_φ measures the number of IID samples required to estimate $\mathbb{E}[\varphi(x)]$ with the same accuracy as $\hat{\varphi}_N$; see e.g. [Geyer \(1992\)](#). To level the computational cost of samplers using multiple gradient steps per iterate, the ESS per gradient evaluation is typically considered as a measure of efficiency. However, the ESS is often intractable and we cannot directly use it to tune the sampler. The Expected Squared Jumping Distance (ESJD) generally serves instead as a proxy ([Pasarica and Gelman, 2010](#)). It is defined as

$$\text{ESJD}_\varphi = \mathbb{E} [(\varphi(X_1) - \varphi(X_0))^2].$$

For the ESJD however, choosing a linear re-scaling is not universal, e.g. [Wang et al. \(2013\)](#) optimizes ESJD/\sqrt{L} for RHMC. This question is discussed in Section 3.4.

For tuning the step-size of HMC, a simple strategy is to target a certain acceptance rate. This strategy is supported by optimal scaling results of [Beskos et al. \(2013\)](#). The tuning of the trajectory length is more challenging however. To overcome this problem, [Hoffman and Gelman \(2014\)](#) proposed a stopping rule that stops trajectories before they turn in on themselves, i.e. committing a ‘‘U-Turn’’. The resultant algorithm (NUTS) requires minimal tuning from the user. NUTS exhibits two limits however: it achieves lower efficiency than optimally tuned RHMC, and the underlying recursion complicates its parallel implementation.

To alleviate these issues, [Hoffman et al. \(2021\)](#); [Sountsov and Hoffman \(2021\)](#) have proposed two adaptive schemes (ChEES and SNAPER) aimed at optimizing the (re-scaled) ESJD with a stochastic optimization routine. The choice of φ is guided by remarking that RHMC’s efficiency is typically bottlenecked by the squared projection on the principal component. These samplers have demonstrated significant gains of efficiency compared to NUTS, especially in the multi-chain context.

Partial momentum refreshment received less attention due to the momentum flips issue. Recently though, GHMC’s improvement of [Neal \(2020\)](#) was leveraged to develop an adaptive sampler named MEADS ([Hoffman and Sountsov, 2022](#)). The damping and the step size were tuned by estimating the maximum eigenvalues of $\nabla(x)$ and $\nabla(\nabla\Phi(x))$. In the sequel, we build upon recent strategies for tuning RHMC to propose tuning heuristics for MALT.

3 TUNING HEURISTICS

In this section, we develop heuristics for tuning each of the following parameters: a diagonal mass matrix M , a damp-

ing γ , a time step h and an integration time τ . We support these heuristics with algorithmic solutions, ensuring that computation and memory costs scale linearly with the dimension of the sampling space.

3.1 Diagonal Preconditioning

The choice of a preconditioner often plays an important role in the efficiency of gradient-based samplers. A common practice consists of tuning M as proportional to the inverse of the (estimated) covariance matrix of Π . However this method does not scale linearly with the dimension. We choose a more lightweight solution by tuning a diagonal mass matrix ([Carpenter et al., 2017](#)) as

$$M = \max(s) \cdot \text{diag}(s)^{-1}$$

where $s = (s_1, \dots, s_d)$ is composed by estimates of the coordinate-wise variances of Π . The mass is normalized¹ by the largest variance to stabilize the learning of the principal component ([Sountsov and Hoffman, 2021](#), Appendix C). In practice, diagonal preconditioning often improves the conditioning of the target distribution without reducing it entirely. Heuristics for the remaining parameters are applied to the pre-conditioned variables $(M^{1/2}x, M^{-1/2}v)$.

3.2 Damping Parameter

The damping is tuned to compensate for the residual conditioning of the target distribution. Choosing the damping requires solving a trade-off between robustness and efficiency. As $\gamma \rightarrow 0$, Algorithm 1 yields HMC with fixed τ , which suffers from poor ergodic properties due to the resonance phenomenon; see Section 2.1. As γ gets large enough, a uniform control of the sampling efficiency becomes possible, typically bottlenecked by the principal component; see [Riou-Durand and Vogrinc \(2022\)](#). Taking γ too large however induces random walk behavior, mitigating the benefits of the kinetic dynamics. Akin to [Hoffman and Sountsov \(2022\)](#), we set the damping² as

$$\gamma = \lambda^{-1/2}$$

where λ is the largest eigenvalue of Σ , which we define as the covariance matrix of $M^{1/2}x$. We estimate the principal component and its eigenvalue with an online PCA method known as CCIPCA ([Weng et al., 2003](#)). This method is characterized by the update

$$w \leftarrow \beta \cdot w + (1 - \beta) \cdot y \cdot (y^\top w) / |w|$$

¹The mass can always be tuned up to a certain constant by rescaling the damping and the time accordingly ([Dalalyan and Riou-Durand, 2020](#), Lemma 1).

²Its optimization is discussed for Gaussian targets in [Riou-Durand and Vogrinc \(2022\)](#), which advocates a slightly underdamped $\gamma = \lambda^{-1/2} < 2\lambda^{-1/2} = \gamma_{\text{critical}}$.

for $y = M^{1/2}(x - m)$ where x is a new sample and m is the mean of Π . The largest eigenvalue and its eigenvector are respectively estimated by

$$\lambda = |w|, \quad z = w/|w|.$$

A formal proof of convergence is established in [Zhang and Weng \(2001\)](#). To mitigate the impact of the warm-up phase, we let $\beta = n/(n + a)$ decay with respect to an amnesia parameter $a > 1$. CCIPCA is known to offer a good trade-off between simplicity and performance. Our adaptive scheme is not constrained by this choice though. It is also compatible with methods such as Oja’s algorithm or IPCA; see [Cardot and Degras \(2018\)](#) for a review.

3.3 Time Step

Guided by optimal scaling results obtained in [Riou-Durand and Vogrinc \(2022\)](#), we propose to tune h by targeting an acceptance rate α^* chosen by the user. The simplicity of this strategy gives MALT a significant advantage compared to GHMC, for which the acceptance rate has to be kept close enough to one to avoid overwhelming the Langevin dynamics with momentum flips ([Bou-Rabee and Vanden-Eijnden, 2010](#)). The acceptance probability of Algorithm 1 unfolds as $\alpha(h) = \mathbb{E}[\exp(-\Delta^+)]$, while we aim to find h such that

$$\alpha(h) - \alpha^* = 0.$$

We leverage the left hand side as a synthetic gradient, for which an unbiased estimator is given by

$$g_h = \exp(-\Delta^+) - \alpha^*.$$

The learning of h results from inputting g_h in a stochastic optimization routine called Adam ([Kingma and Ba, 2014](#)); see Algorithm 2 (gradient ascent).

Algorithm 2 Adaptive moment estimation (Adam)

Hyperparameters: $\delta, \beta_1, \beta_2, \epsilon$

- 1: **function** ADAM($\theta, m, s, g; n$)
 - 2: Set $m' \leftarrow \beta_1 m + (1 - \beta_1)g$ and $\hat{m} \leftarrow m'/(1 - \beta_1^n)$
 - 3: Set $s' \leftarrow \beta_2 s + (1 - \beta_2)g^2$ and $\hat{s} \leftarrow s'/(1 - \beta_2^n)$
 - 4: Update $\theta' \leftarrow \theta + \delta \hat{m}/(\sqrt{\hat{s}} + \epsilon)$
 - 5: **return** θ', m', s'
 - 6: **end function**
-

3.4 Integration Time

We propose to tune τ by focusing on the sampling efficiency of the principal component z , measured by ESS_φ for $\varphi(x) = (z^\top M^{1/2}(x - m))^2$. Akin to [Sountsov and Hoffman \(2021\)](#), we consider the variance of the projection rather than its mean to avoid encouraging antithetic behaviors. The computational cost of Algorithm 1 is proportional

to the number of gradient evaluations $L = \lceil \tau/h \rceil$, thus increasing τ is profitable only if the ESS increases at least linearly. A natural criterion to optimize is therefore

$$\mathcal{C}(\tau) \triangleq \text{ESS}_\varphi(\tau)/\tau.$$

However, deriving an unbiased estimator of the ESS or its gradient is challenging. A common practice is to instead work with the ESJD, which is more suitable for use in stochastic optimization routines such as Adam. We remark that the choice of a linear re-scaling for the ESJD is not unanimous in the literature. For instance a re-scaling by $\tau^{1/2}$ was chosen by [Wang et al. \(2013\)](#) for RHMC. In between these two options, we consider a family of criterions indexed by $\rho \in [0, 1]$, defined as

$$\mathcal{C}_\rho(\tau) \triangleq \text{ESJD}_\varphi(\tau)/\tau^{(1+\rho)/2}.$$

The choice $\rho = 1$ is supported by the fact that $\text{ESJD}_\varphi(\tau)$ is proportional to an upper bound on $\text{ESS}_\varphi(\tau)$ whenever $\rho_\varphi(\tau) \triangleq \text{Corr}(\varphi(X_\tau), \varphi(X_0)) \geq 0$ ([Sountsov and Hoffman, 2021](#)). MALT being Π -reversible, we obtain from [Geyer \(1992\)](#) and [Sokal \(1997\)](#) that $\mathcal{C}(\tau)$ is upper bounded by the criterion

$$\bar{\mathcal{C}}(\tau) \triangleq \frac{N}{\tau} \cdot \frac{1 - \rho_\varphi(\tau)}{1 + \rho_\varphi(\tau)}.$$

Indeed, if $\rho_\varphi(\tau) \geq 0$ then $\bar{\mathcal{C}}(\tau) \geq N(1 - \rho_\varphi(\tau))/\tau$, and $\text{ESJD}_\varphi(\tau) \propto N(1 - \rho_\varphi(\tau))$ implies that $\mathcal{C}_1(\tau)$ is proportional to an upper bound on $\bar{\mathcal{C}}(\tau)$. In this work, we derive a method to target directly

$$\tau^* = \underset{\tau > 0}{\text{argmax}} \bar{\mathcal{C}}(\tau)$$

by noticing that τ^* coincides with the optimizer of \mathcal{C}_ρ when $\rho = \rho_\varphi(\tau^*)$. We therefore propose to target \mathcal{C}_ρ with Algorithm 2 while tuning ρ adaptively as an estimate of the first autocorrelation of φ (this heuristic is further supported by a fixed point argument in the supplement).

Differentiating with respect to τ yields $\mathcal{C}'_\rho(\tau) = 0$ if and only if

$$\frac{d}{d\tau} \text{ESJD}_\varphi(\tau) - \frac{1 + \rho}{2\tau} \cdot \text{ESJD}_\varphi(\tau) = 0.$$

Considering the solution of the SDE (3) at stationarity, $(\varphi(x_\tau) - \varphi(x_0))^2$ is an unbiased estimator of $\text{ESJD}_\varphi(\tau)$. In addition, since $dx_\tau = M^{-1}v_\tau d\tau$, an unbiased estimator of the derivative of $\text{ESJD}_\varphi(\tau)$ is given by

$$\delta(x_\tau, x_0, v_\tau) \triangleq 2(\nabla\varphi(x_\tau)^\top M^{-1}v_\tau) \cdot (\varphi(x_\tau) - \varphi(x_0))$$

where $\nabla\varphi(x) = 2(z^\top M^{1/2}(x - m)) \cdot M^{1/2}z$. We leverage the time reversibility of (3) and remark that (x_τ, x_0, v_τ) and $(x_0, x_\tau, -v_0)$ have the same distribution. This enables the use of a reduced variance estimator, defined as

$$g(x_\tau, v_\tau, x_0, v_0) \triangleq \frac{1}{2} \left(\delta(x_\tau, x_0, v_\tau) + \delta(x_0, x_\tau, -v_0) \right).$$

The two components of this averaged estimator are asymptotically uncorrelated as $\tau \rightarrow \infty$ (see supplement). Therefore, considering g in place of δ is expected to reduce the variance by a factor 2 as τ gets large. This variance reduction method is applied to both MALT and RHMC in the experiments of Section 5, to speed up the adaptation. As an example of this effect for MALT on the Hierarchical Linear Model, the variance of the τ estimate goes from 0.11 to 0.057 when the improved estimator is used.

4 ADAPTATION WITH MANY CHAINS

We now describe an Adaptive MCMC algorithm for tuning the hyperparameters of MALT which combines the heuristics described in the previous section. We assume that we can run K chains in parallel, which allows us to share statistical strength between them as well as being amenable to modern SIMD hardware accelerators. The heuristics discussed previously are readily adapted to this setting. Algorithm 3 outlines the resultant *Adaptive MALT*. The moments m, s, w, m_2, s_2 required by the optimization of τ and the tuning of M and γ are updated with respect to the functions $f_\psi(x; \theta)$ summarized in Table 1. Figure 1 shows an example of the adaptation trace produced by Adaptive MALT.

Algorithm 3 Adaptive MALT

Updates: Parallel states: X_1, \dots, X_K . Online estimates of $\theta = (m, s, w, m_2, s_2)$ and $h, m_h, s_h, \tau, m_\tau, s_\tau, c$.

Hyperparameters: $N_{\text{adapt}}, N_{\text{clip}}, a$

- 1: **for** $n = 1$ to N_{adapt} **do**
 - 2: \triangleright Compute the mass matrix and friction.
 - 3: set $M \leftarrow \max(s) \cdot \text{diag}(s)^{-1}$ and $\gamma \leftarrow |w|^{-1/2}$
 - 4: **if** $n \leq N_{\text{clip}}$ **then** set $\tau \leftarrow h$
 - 5: **end if**
 - 6: set $\beta \leftarrow n/(n+a)$ and $\rho \leftarrow c^+/s_2$
 - 7: \triangleright Perform MALT updates.
 - 8: **for** $k = 1$ to K **do in parallel**
 - 9: $X_k, v_\tau, x_0, v_0, \Delta_k \leftarrow \text{MALT}(X_k; M, \gamma, h, \tau)$
 - 10: $g_k \leftarrow g(X_k, v_\tau, x_0, v_0) - \frac{1+\rho}{2\tau}(\varphi(X_k) - \varphi(x_0))^2$
 - 11: $c_k \leftarrow (\varphi(X_k) - m_2)(\varphi(x_0) - m_2)$
 - 12: **end for**
 - 13: \triangleright Adapt step size.
 - 14: $g_h \leftarrow \frac{1}{K} \sum_k e^{-\Delta_k^+} - \alpha^*$
 - 15: $\log(h), m_h, s_h \leftarrow \text{ADAM}(\log(h), m_h, s_h, g_h; n)$
 - 16: \triangleright Adapt trajectory length.
 - 17: $g_\tau \leftarrow \frac{1}{K} \sum_k g_k$
 - 18: $\log(\tau), m_\tau, s_\tau \leftarrow \text{ADAM}(\log(\tau), m_\tau, s_\tau, g_\tau; n)$
 - 19: \triangleright Adapt online estimates.
 - 20: $c \leftarrow \beta c + (1 - \beta) \frac{1}{K} \sum_k c_k$
 - 21: **for** $\psi \in \{m, s, w, m_2, s_2\}$ **do**
 - 22: $\psi \leftarrow \beta \psi + (1 - \beta) \frac{1}{K} \sum_k f_\psi(X_k; \theta)$
 - 23: **end for**
 - 24: **end for**
-

Table 1: Summary of the updates for $y_x = M^{1/2}(x - m)$, $M = \max(s) \cdot \text{diag}(s)^{-1}$ and $z = w/|w|$ in Algorithm 3.

Parameter	ψ	$f_\psi(x; \theta)$
Vector of means	m	x
Vector of variances	s	$(x - m)^{\odot 2}$
Principal eigenvector	w	$(z^\top y_x) \cdot y_x$
Mean squared projection	m_2	$(z^\top y_x)^2$
Variance squared projection	s_2	$((z^\top y_x)^2 - m_2)^2$




Figure 1: Trace of the τ and γ hyperparameters from one run of Adaptive MALT on the Hierarchical Linear Model; see Section 5. Every dot represents 5 steps of the algorithm. In the background, the surface plot is the minimum (across dimensions) ESS/grad for estimating the centered second moment, computed via grid search.

We consider here a finite adaptation scheme, for which ergodicity is ensured by the ergodicity of each MALT kernel; see [Roberts and Rosenthal \(2007, Proposition 2\)](#). In Algorithm 3, the time step and the integration time are optimized by inputting the averaged gradient estimates g_h and g_τ into Adam, defined in Algorithm 2. The τ optimization targets the criterion C_ρ by setting ρ adaptively as described in Section 3.4. This adaptation can be turned off by setting $\rho = 1$, in which case the estimation of m_2 and s_2 is no longer required; see Section 5 for a comparison. The updates we present are averaged across chains using an arithmetic mean to preserve unbiasedness for g_h and g_τ . For robustness, a common practice is to use instead the harmonic mean ([Hoffman et al., 2021](#)) to average the acceptance rates across chains when tuning the step size.

While we frame the Adaptive MALT algorithm for the case where $K > 1$, it is defined to remain functional even in cases where $K = 1$. This is a setting where there is insuffi-

cient memory to run multiple chains in parallel, or a hardware accelerator is not available, or already well-utilized by SIMD computation within Φ and $\nabla\Phi$. We observe that setting $K = 1$ requires lower learning rates and longer chains to allow for the hyperparameters to converge.

It is interesting to compare the implementation of many-chain MALT and RHMC. By construction, MALT uses a single integration time τ while RHMC needs to randomize it every iteration. A straightforward implementation of many chain RHMC would require each iteration to simulate the number of leapfrog steps necessary for the chain that has the longest integration time draw, causing wasted computation for the rest of the chains. Practical implementation of RHMC (Hoffman et al., 2021; Sountsov and Hoffman, 2021) instead choose to share the random τ draw across chains, trading efficiency for correlation across chains. In contrast, many-chain MALT does not need to make similar trade-offs.

When comparing MALT and GHMC as implemented by MEADS, one key difference is the estimator of largest eigenvalue λ . MALT, as discussed in Section 3.2 uses CCIPCA which has compute complexity of $O(Kd)$. MEADS, due to its non-adaptive nature, requires the use of an estimator that can return good estimates of λ using only the current state of the MCMC chain. MEADS uses an estimator based on approximating the ratio $\text{tr}(\Sigma^2)/\text{tr}(\Sigma) \approx \lambda$ with an algorithm with a compute complexity of $O(K^2d)$. For large number of chains, the additional computational burden can be significant, which is compounded by the fact that MEADS requires a large number of chains to provide accurate estimates whereas Adaptive MALT does not.

5 EXPERIMENTS

We validate Adaptive MALT and compare it to a number of baselines by targeting the posterior distributions of several Bayesian models. All the models are sourced from the Inference Gym (Sountsov et al., 2020), and are summarized in Table 2. In addition to Adaptive MALT, we also examine the performance of RHMC (as implemented by SNAPER; Sountsov and Hoffman (2021)), GHMC (as implemented by MEADS; Hoffman and Sountsov (2022)) and NUTS (following the implementation from Sountsov and Hoffman (2021)). For MALT we examine the effect of tuning ρ or keeping it fixed at 1 (see Section 3.4). We use the performance of the latter configuration to normalize the results from the remaining algorithms. For RHMC we sample the trajectory lengths from a uniform distribution (*RHMC (Unif)*) or an exponential distribution (*RHMC (Exp)*), both with mean τ . Additionally, we use the improved estimator for computing the trajectory criterion gradient as described in Section 3.4. The target acceptance rate is chosen as $\alpha^* = 80\%$, slightly above the asymptotically optimal 65% rate; see Beskos et al. (2013) for HMC

and Riou-Durand and Vogrinc (2022) for MALT. This is a common practice to ensure robustness with minimal loss of efficiency for HMC (Betancourt et al., 2014). A detailed study of this tradeoff for MALT is beyond our scope.

We run all algorithms for 128 parallel chains, with 5000 adaptive warmup iterations, followed by 400 non-adaptive warmup iterations, and finally 1600 sampling iterations used for the collection of statistics. We repeat each simulation run 20 times with different PRNG seeds. All algorithms are implemented using the Python package JAX (Bradbury et al., 2018) and run on a GPU³.

The number of chains is chosen large enough to meet MEADS’ requirements for its parameter tuning scheme. The MEADS implementation is also sensitive to initialization (in the original work, a MAP estimate is used), so to level the playing field we initialize all algorithms with samples derived from Automatic Differentiation Variational Inference (ADVI) (Kucukelbir et al., 2016).

We report the minimum Effective Sample Size (ESS) of the sampling phase among the d squared components. The latter is relevant for estimating the marginal variances of the posterior distribution. Notably, this metric is independent from the estimation of the principal component, and accounts for resonances (see Section 3.2). The ESS is normalized either by the number of $\nabla\Phi$ evaluations (ESS/grad), or the number of MCMC iterations (ESS/iter). The first metric reflects the efficiency per unit of computation, as gradient evaluations typically dominate the computation time in high dimension. The second metric reflects the efficiency per unit of storage, focusing on situations when memory is at a premium, and many iterations must be stored for analysis. Such a situation can arise when fitting a high-dimensional model using a large number of chains.

There is a subtlety when computing ESS/grad for NUTS in a parallel-chain setting, as its ragged computation will necessarily waste some computation compared to the other algorithms. To address this we compute the number of gradients as both the maximum across chains for each step (reflecting the actual implementation) and the mean across chains (labeled as *NUTS (avg)*). The latter measures performance that is closer to the non-SIMD single chain implementation, but is difficult to reach in a SIMD setting. For MEADS, we follow Hoffman and Sountsov (2022) and thin the chain by a factor of 10, meaning that each iteration consumes 10 gradient evaluations. Both metrics ESS/grad and ESS/iter are sensitive to the choice of thinning, but tuning it optimally is outside the scope of this work.


The results are summarized in Figure 2.

When considering ESS/grad, Adaptive MALT is compet-

³Implementation of the experiments is available at https://github.com/tensorflow/probability/tree/main/discussion/adaptive_malt

Table 2: Bayesian models used for the experiments. The code snippets refer to the Python class invocation to construct the model (the classes are all found in the `inference_gym.targets` module).

Model	d	Inference Gym Snippet
Logistic Regression	25	<code>GermanCreditNumericLogisticRegression()</code>
Brownian Bridge	32	<code>BrownianMotionUnknownScalesMissingMiddleObservations(use_markov_chain=True)</code>
Sparse Logistic Regression	51	<code>GermanCreditNumericSparseLogisticRegression(positive_constraint_fn="softplus")</code>
Hierarchical Linear Model	97	<code>RadonContextualEffectsIndiana()</code>
Item Response Theory	501	<code>SyntheticItemResponseTheory()</code>
Stochastic Volatility	2519	<code>VectorizedStochasticVolatilitySP500()</code>


 Figure 2: Minimum ESS/grad and ESS/iter for estimating the centered second moment across algorithms and models. All values are normalized by the corresponding efficiency measure attained by Adaptive MALT with fixed $\rho = 1$. Values are 10th percentile across 20 runs. Errors bars are 3 SEM estimated via bootstrap, but are often too small to see.

itive with the best algorithms. It outperforms MEADS (GHMC) in 5 out of the 6 problems. It outperforms NUTS in all 6 problems. However, it is often slightly less efficient than RHMC with uniform jitter. Aside from examining MALT’s performance, our experiments also highlight one hitherto under-explored design choice of RHMC: the distribution of trajectory lengths. We observe that the uniform distribution produces a higher efficiency sampler than the exponential distribution, which has been previously studied for theoretical convenience (Riou-Durand and Vogrinc, 2022; Bou-Rabee and Sanz-Serna, 2017). A possible explanation for this is that integration times sampled from an exponential distribution which exceed 2τ are wasteful due to performing a U-turn. Compared to setting $\rho = 1$, we observe that targeting $\rho = \rho_\phi(\tau^*)$ only improves the efficiency in one problem: the Item Response Theory model. Despite the fact that adapting ρ targets a sharper upper bound on the ESS (in continuous time), it also induces a smaller step-size as the trajectory length increases. As a result, the numerical efficiency is optimized by choosing slightly shorter trajectories.

When looking at ESS/iter, Adaptive MALT performs better than the other algorithms in most problems. This is

mainly due to the fact that MALT prefers slightly longer trajectories than RHMC. This is emphasized when setting ρ adaptively. Comparing both algorithms based on Langevin dynamics, we see that MALT often outperforms MEADS (GHMC), all the while not requiring choosing a thinning schedule and being amenable to few-chain adaptation.

6 DISCUSSION

This paper proposes tuning guidelines for each of MALT’s hyperparameters. The tuning of MALT leverages recent strategies for tuning RHMC (Hoffman et al., 2021; Soutsov and Hoffman, 2021), consolidated with: (i) the use of a parameter-free online PCA method (CCIPCA), to estimate the principal component and its eigenvalue, (ii) an adaptive re-scaling of the ESJD which targets a sharper upper bound on the ESS, (iii) a reduced variance estimator of the gradient of τ . These methods are new and applicable beyond the scope of MALT, e.g. (iii) enables a faster adaptation of RHMC in Section 5.

The resulting Adaptive MALT is user-friendly, readily applicable to a broad range of targets and benefits from a high-performance implementation in JAX. Unlike NUTS,

it is amiable to hardware accelerators and can be run efficiently using many chains in parallel. The benefits of running many chains include better adaptation (as discussed in this paper and references) but also opportunities to make sampling phases arbitrarily short and construct more principled convergence diagnostics (Margossian et al., 2022). With the many chains regime becoming more standard, MALT’s effective use of memory is particularly compelling.

Compared to GHMC, the MALT algorithm bypasses the momentum-flips problem by applying the Metropolis correction to entire Langevin trajectories. Our work validates this qualitative advantage by showing that Adaptive MALT outperforms GHMC numerically, even when improved by a slice-sampler (Neal, 2020; Hoffman and Sountsov, 2022). MALT’s ability to avoid momentum flips paves the way for developing new efficient algorithms, e.g. based on other types of partial momentum refreshment.

Acknowledgments

We thank Matthew D. Hoffman as well as the anonymous reviewers for their helpful comments and suggestions. LRD was supported by the EPSRC grant EP/R034710/1. JV was supported by the EPSRC grant EP/T004134/1. SP was supported by the EPSRC grant EP/R018561/1.

References

- Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M., and Stuart, A. (2013). Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534.
- Betancourt, M., Byrne, S., and Girolami, M. (2014). Optimizing the integrator step size for Hamiltonian Monte Carlo. *arXiv preprint arXiv:1411.6669*.
- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., and Goodman, N. D. (2018). Pyro: Deep universal probabilistic programming.
- Bou-Rabee, N. and Sanz-Serna, J. M. (2017). Randomized Hamiltonian Monte Carlo. *Ann. Appl. Probab.*, 27(4):2159–2194.
- Bou-Rabee, N. and Vanden-Eijnden, E. (2010). Pathwise accuracy and ergodicity of metropolized integrators for SDEs. *Comm. Pure Appl. Math.*, 63(5):655–696.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs.
- Campos, C. M. and Sanz-Serna, J. M. (2015). Extra chance generalized hybrid Monte Carlo. *J. Comput. Phys.*, 281:365–374.
- Cao, Y., Lu, J., and Wang, L. (2019). On explicit l^2 -convergence rate estimate for underdamped langevin dynamics. *arXiv preprint arXiv:1908.04746*.
- Cardot, H. and Degras, D. (2018). Online principal component analysis in high dimension: Which algorithm to choose? *International Statistical Review*, 86(1):29–50.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, 76(1):1–32.
- Dalalyan, A. S. and Riou-Durand, L. (2020). On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli*, 26(3):1956–1988.
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216–222.
- Eberle, A., Guillin, A., and Zimmer, R. (2019). Couplings and quantitative contraction rates for Langevin dynamics. *Ann. Probab.*, 47(4):1982–2010.
- Geyer, C. J. (1992). Practical markov chain monte carlo. *Statistical science*, pages 473–483.
- Hoffman, M., Radul, A., and Sountsov, P. (2021). An adaptive-mcmc scheme for setting trajectory lengths in hamiltonian monte carlo. In *International Conference on Artificial Intelligence and Statistics*, pages 3907–3915. PMLR.
- Hoffman, M. D. and Gelman, A. (2014). The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1593–1623.
- Hoffman, M. D. and Sountsov, P. (2022). Tuning-free generalized hamiltonian monte carlo. In *International Conference on Artificial Intelligence and Statistics*, pages 7799–7813. PMLR.
- Horowitz, A. M. (1991). A generalized guided monte carlo algorithm. *Physics Letters B*, 268(2):247–252.
- Kennedy, A. D. and Pendleton, B. (2001). Cost of the generalised hybrid Monte Carlo algorithm for free field theory. *Nuclear Phys. B*, 607(3):456–510.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. M. (2016). Automatic differentiation variational inference.
- Lao, J. and Dillon, J. (2019). Unrolled implementation of no-u-turn sampler. https://github.com/tensorflow/probability/blob/master/discussion/technical_note_on_unrolled_nuts.md.

- Lao, J., Suter, C., Langmore, I., Chimisov, C., Saxena, A., Sountsov, P., Moore, D., Saurous, R. A., Hoffman, M. D., and Dillon, J. V. (2020). tfp.mcmc: Modern markov chain monte carlo tools built for modern hardware.
- Margossian, C. C., Hoffman, M. D., Sountsov, P., Riou-Durand, L., Vehtari, A., and Gelman, A. (2022). Nested \hat{R} : Assessing the convergence of markov chains monte carlo when running many short chains. *arXiv:2110.13017*.
- Mira, A. (2001). On Metropolis-Hastings algorithms with delayed rejection. *Metron*, 59(3-4):231–241.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. pages 113–162.
- Neal, R. M. (2020). Non-reversibly updating a uniform $[0, 1]$ value for metropolis accept/reject decisions. *arXiv preprint arXiv:2001.11950*.
- Park, J. and Atchadé, Y. (2020). Markov chain Monte Carlo algorithms with sequential proposals. *Stat. Comput.*, 30(5):1325–1345.
- Pasarica, C. and Gelman, A. (2010). Adaptively scaling the metropolis algorithm using expected squared jumped distance. *Statist. Sinica*, 20(1):343–364.
- Phan, D. and Pradhan, N. (2019). Iterative NUTS. <https://github.com/pyro-ppl/numpyro/wiki/Iterative-NUTS>.
- Radul, A., Patton, B., Maclaurin, D., Hoffman, M., and A Saurous, R. (2020). Automatically batching control-intensive programs for modern accelerators. *Proceedings of Machine Learning and Systems*, 2:390–399.
- Riou-Durand, L. and Vogrinc, J. (2022). Metropolis adjusted langevin trajectories: a robust alternative to hamiltonian monte carlo. *arXiv preprint arXiv:2202.13230*.
- Roberts, G. O. and Rosenthal, J. S. (2007). Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *Journal of Applied Probability*, 44(2):458–475.
- Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. (2016). Probabilistic programming in python using PyMC3. *PeerJ Computer Science*, 2:e55.
- Sokal, A. (1997). Monte carlo methods in statistical mechanics: foundations and new algorithms. In *Functional integration*, pages 131–192. Springer.
- Sountsov, P. and Hoffman, M. D. (2021). Focusing on difficult directions for learning hmc trajectory lengths. *arXiv preprint arXiv:2110.11576*.
- Sountsov, P., Radul, A., and contributors (2020). Inference gym. https://pypi.org/project/inference_gym.
- Wang, Z., Mohamed, S., and Freitas, N. (2013). Adaptive hamiltonian and riemann manifold monte carlo. In *International conference on machine learning*, pages 1462–1470. PMLR.
- Weng, J., Zhang, Y., and Hwang, W.-S. (2003). Candid covariance-free incremental principal component analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):1034–1040.
- Zhang, Y. and Weng, J. (2001). Convergence analysis of complementary candid incremental principal component analysis. *Michigan State University*.

Supplementary Material of “Adaptive Tuning for Metropolis Adjusted Langevin Trajectories”

Appendix A ADAPTIVE REGULARIZATION OF THE ESJD

In this section, we derive the adaptive regularization of the ESJD presented in Section 3.4. We note $(x_\tau, v_\tau)_{\tau \geq 0}$ the stationary solution of the Langevin SDE defined in Section 2.2. We consider the criterions $\bar{\mathcal{C}}(\tau)$ and $\mathcal{C}_\rho(\tau)$ for $\rho \in [0, 1]$, defined in Section 3.4. Our analysis holds for any continuously differentiable function $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$ such that each expectation of this section exists. We make repeated use of the short-hand notations

$$J(\tau) \triangleq \mathbb{E}[(\varphi(X_\tau) - \varphi(X_0))^2], \quad \rho(\tau) \triangleq \text{Corr}(\varphi(X_\tau), \varphi(X_0)).$$

We highlight that $\rho \in [0, 1]$ denotes a tuning parameter, while $\rho(\tau)$ stands for the auto-correlation function after time τ . Since $J(\tau) \propto 1 - \rho(\tau)$, we get

$$\bar{\mathcal{C}}(\tau) \propto \frac{J(\tau)}{\tau(1 + \rho(\tau))}, \quad \mathcal{C}_\rho(\tau) = \frac{J(\tau)}{\tau^{(1+\rho)/2}}.$$

Assumption 1. *The criterions $\bar{\mathcal{C}}$ and \mathcal{C}_ρ have unique maximizers for every $\rho \in [0, 1]$, denoted*

$$\tau^* \triangleq \operatorname{argmax}_{\tau > 0} \bar{\mathcal{C}}(\tau), \quad \tau_\rho^* \triangleq \operatorname{argmax}_{\tau > 0} \mathcal{C}_\rho(\tau).$$

Assumption 2. *The map $\tau \mapsto \rho(\tau)$ is positive decreasing on $[\tau_1^*, \tau_0^*]$.*

On Gaussian targets, these two conditions are satisfied for the map $\varphi(x) = (z^\top M^{1/2}(x - m))^2$ considered in Section 3.4. In such a case, the criterions $\bar{\mathcal{C}}(\tau)$ and $\mathcal{C}(\tau) = \text{ESS}_\varphi(\tau)/\tau$ also coincide, due to the geometric decay of the autocorrelations.

For any $\tau > 0$, the auto-correlation $\rho(\tau)$ can be approximated with a long run of MALT. Conversely for any $\rho \in [0, 1]$, the optimizer of \mathcal{C}_ρ can be approximated from a long run of Adam. In Section 3.4, we propose to target adaptively \mathcal{C}_ρ with Adam while setting ρ equal to the auto-correlation of the MALT chain. We support this strategy with a fixed point argument in the next proposition.

Proposition 1. *Suppose Assumptions 1 and 2 hold. Define $f(s) \triangleq \tau_{\rho(s)}^*$ for any $s \geq 0$ such that $\rho(s) \geq 0$. Then τ^* is the unique attracting fixed point of f . Furthermore, the sequence $s_{k+1} = f(s_k)$ converges to τ^* for any start $s_0 \in [\tau_1^*, \tau_0^*]$.*

The claim of Proposition 1 is established by proving the following properties:

- (i) τ^* is the unique fixed point of f .
- (ii) $\rho \mapsto \tau_\rho^*$ is continuously decreasing on $[0, 1]$.
- (iii) f is continuously increasing on $[\tau_1^*, \tau_0^*]$.
- (iv) $\forall s_0 \in [\tau_1^*, \tau_0^*]$ the sequence $s_{k+1} = f(s_k)$ converges to τ^* .

Proof of (i). Define the synthetic gradients \bar{g} and g_ρ as

$$\bar{g}(\tau) \triangleq J'(\tau) - \frac{1 + \rho(\tau)}{2\tau} J(\tau), \quad g_\rho(\tau) \triangleq J'(\tau) - \frac{1 + \rho}{2\tau} J(\tau).$$

We remark that $\bar{\mathcal{C}}'(\tau) = 0 \Leftrightarrow \bar{g}(\tau) = 0$, and that $\mathcal{C}'_\rho(\tau) = 0 \Leftrightarrow g_\rho(\tau) = 0$, where the first equivalence follows from the identity $J'(\tau)(1 - \rho(\tau)) = -J(\tau)\rho'(\tau)$. As a consequence, τ^* and τ_ρ^* are the unique zeros of \bar{g} and g_ρ respectively, by Assumption 1. Finally we obtain that $\tau = f(\tau) \Leftrightarrow \tau = \tau^*$, because $g_{\rho(\tau)}$ coincides with \bar{g} and τ^* is its unique zero.

Proof of (ii). Let $\rho \in [0, 1]$. From Assumption 1, we deduce that the map $\psi(s) \triangleq 2sJ'(s)/J(s) - 1$ is continuously decreasing in a neighborhood of τ_ρ^* . This observation follows from the expression of g_ρ defined above. As a consequence, ψ is locally invertible and the map $u \mapsto \psi^{-1}(u) = \tau_u^*$ is continuously decreasing in a neighborhood of ρ .

Proof of (iii). The claim follows from noticing that $f : s \mapsto \psi^{-1}(\rho(s))$ is a composition of two continuously decreasing functions, from (ii) and Assumption 2.

Proof of (iv). We first show that $\tau_1^* < f(\tau_1^*) < \tau^* < f(\tau_0^*) < \tau_0^*$. Indeed, since $\tau^* = \tau_{\rho(\tau^*)}$ and $\rho(\tau^*) \in (0, 1)$, we obtain from (ii) that $\tau_1^* < \tau^* < \tau_0^*$. Using Assumption 2, we deduce that $1 > \rho(\tau_1^*) > \rho(\tau^*) > \rho(\tau_0^*) > 0$. We conclude by applying (ii) to these last inequalities.

From (i) and (iii), we know that τ^* is the unique fixed point of the continuous map f . Therefore, a similar inequality holds for any intermediate points: for any $s, s' \in [\tau_1^*, \tau_0^*]$ such that $s < \tau^* < s'$, we have $s < f(s) < \tau^* < f(s') < s'$.

Now, let $s_0 \in [\tau_1^*, \tau^*)$ and $s_{k+1} = f(s_k)$. We obtain from (i) and (iii) that $s_k < f(s_k) = s_{k+1} < f(\tau^*) = \tau^*$. The sequence (s_k) is increasing and upper bounded, thus it is convergent. The claim follows from (i) since the limit is necessarily a fixed point of f . The case $s_0 \in (\tau^*, \tau_1^*]$ is symmetric.

Appendix B VARIANCE REDUCTION

In this section, we derive a quantification of the variance reduction proposed in Section 3.4, asymptotically as $\tau \rightarrow +\infty$. We also illustrate numerically the variance reduction on the Hierarchical Linear Model; see Section 5.

B.1 Asymptotic Quantification

Similarly to the previous section, we note $(x_\tau, v_\tau)_{\tau \geq 0}$ the stationary solution of the Langevin SDE defined in Section 2.2. We also note that our analysis holds for any continuously differentiable function $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$ such that each expectation of this section exists.

From Section 3.4, we recall the unbiased estimators of the derivative of ESJD $_\varphi$, defined as

$$\begin{aligned} \delta(x_\tau, x_0, v_\tau) &= (\varphi(x_\tau) - \varphi(x_0)) \nabla \varphi(x_\tau)^\top M^{-1} v_\tau, \\ g(x_\tau, v_\tau, x_0, v_0) &= \frac{1}{2} \left(\delta(x_\tau, x_0, v_\tau) + \delta(x_0, x_\tau, -v_0) \right). \end{aligned}$$

Assumption 3. For every $\psi, \omega \in \mathbb{L}_2(\Pi^*)$, we have $\mathbb{E}[\psi(x_\tau, v_\tau) \omega(x_0, v_0)] \rightarrow \mathbb{E}[\psi(x_0, v_0)] \cdot \mathbb{E}[\omega(x_0, v_0)]$ as $\tau \rightarrow +\infty$.

This assumption holds whenever the kinetic Langevin diffusion mixes in \mathbb{L}_2 ; see Section 2.2 for references. In the next proposition, we show that choosing the estimator g instead of δ leads to a variance reduction of a factor $1/2$ as $\tau \rightarrow +\infty$. Studying the speed of convergence towards this factor is beyond our work, although we expect that an exponential decay can be derived under stronger assumptions, provided that these ensure geometric \mathbb{L}_2 -mixing of the Langevin diffusion.

Proposition 2. Suppose Assumption 3 holds. Then $\delta(x_\tau, x_0, v_\tau)$ and $\delta(x_0, x_\tau, -v_0)$ are asymptotically uncorrelated as $\tau \rightarrow +\infty$. Therefore

$$\lim_{\tau \rightarrow +\infty} \frac{\mathbb{V}(g(x_\tau, v_\tau, x_0, v_0))}{\mathbb{V}(\delta(x_\tau, x_0, v_\tau))} = \frac{1}{2}.$$

Since the estimators $\delta(x_\tau, x_0, v_\tau)$ and $\delta(x_0, x_\tau, -v_0)$ have the same distribution, we remark that

$$\mathbb{V}(g(x_\tau, v_\tau, x_0, v_0)) = \frac{1}{2} \mathbb{V}(\delta(x_\tau, x_0, v_\tau)) + \frac{1}{2} \text{Cov}(\delta(x_\tau, x_0, v_\tau), \delta(x_0, x_\tau, -v_0)).$$

The claim of Proposition 2 follows by showing that the second term vanishes. We prove that, as $\tau \rightarrow +\infty$, we have:

- (i) $\mathbb{E}[\delta(x_\tau, x_0, v_\tau)] \rightarrow 0$.
- (ii) $\mathbb{E}[\delta(x_\tau, x_0, v_\tau) \delta(x_0, x_\tau, -v_0)] \rightarrow 0$.

Proof of (i). We make repeated use of the fact that v_0 is centered and independent from x_0 . As a consequence, for every Π -integrable function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$, we have $\mathbb{E}[f(x_0)^\top M^{-1}v_0] = \mathbb{E}[f(x_0)]^\top M^{-1}\mathbb{E}[v_0] = 0$. We combine this property with Assumption 3, and obtain that

$$\begin{aligned} \lim_{\tau \rightarrow +\infty} \mathbb{E}[\delta(x_\tau, x_0, v_\tau)] &= \lim_{\tau \rightarrow +\infty} \left(\mathbb{E}[\varphi(x_\tau) \nabla \varphi(x_\tau)^\top M^{-1}v_\tau] - \mathbb{E}[\varphi(x_0) \nabla \varphi(x_0)^\top M^{-1}v_0] \right) \\ &= \mathbb{E}[(\varphi(x_0) \nabla \varphi(x_0))^\top M^{-1}v_0] - \mathbb{E}[\varphi(x_0)] \cdot \mathbb{E}[\nabla \varphi(x_0)^\top M^{-1}v_0] = 0. \end{aligned}$$

Proof of (ii). The time-reversibility of the Langevin dynamics implies that $(x_\tau, v_\tau, x_0, v_0)$ and $(x_0, -v_0, x_\tau, -v_\tau)$ have the same distribution; see Section 3.4. Therefore

$$\begin{aligned} \mathbb{E}[\delta(x_\tau, x_0, v_\tau) \delta(x_0, x_\tau, -v_0)] &= \mathbb{E}[(\varphi(x_\tau) - \varphi(x_0))^2 (\nabla \varphi(x_\tau)^\top M^{-1}v_\tau) (\nabla \varphi(x_0)^\top M^{-1}v_0)] \\ &= 2 \mathbb{E}[\varphi(x_\tau)^2 (\nabla \varphi(x_\tau)^\top M^{-1}v_\tau) (\nabla \varphi(x_0)^\top M^{-1}v_0)] \\ &\quad + 2 \mathbb{E}[\varphi(x_\tau)\varphi(x_0) (\nabla \varphi(x_\tau)^\top M^{-1}v_\tau) (\nabla \varphi(x_0)^\top M^{-1}v_0)]. \end{aligned}$$

Again, we combine Assumption 3 with the fact that v_0 is centered and independent from x_0 . We obtain that the two last terms vanish as $\tau \rightarrow +\infty$, since

$$\begin{aligned} \mathbb{E}[\varphi(x_\tau)^2 (\nabla \varphi(x_\tau)^\top M^{-1}v_\tau) (\nabla \varphi(x_0)^\top M^{-1}v_0)] &\rightarrow \mathbb{E}[\varphi(x_0)^2 (\nabla \varphi(x_0)^\top M^{-1}v_0)] \cdot \mathbb{E}[\nabla \varphi(x_0)^\top M^{-1}v_0] = 0, \\ \mathbb{E}[\varphi(x_\tau)\varphi(x_0) (\nabla \varphi(x_\tau)^\top M^{-1}v_\tau) (\nabla \varphi(x_0)^\top M^{-1}v_0)] &\rightarrow \left(\mathbb{E}[\varphi(x_0) (\nabla \varphi(x_0)^\top M^{-1}v_0)] \right)^2 = 0. \end{aligned}$$

B.2 Numerical Illustration

An example trace of τ over time using the estimators δ (fwd), and g (fwd+rev), can be seen in Figure 3. The variation in the trace is caused primarily by the noise in the gradient estimate, and is reduced approximately by 2 when the improved estimator is used.




Figure 3: Trace of the τ hyperparameter near convergence from two runs of Adaptive MALT on the Hierarchical Linear Model; see Section 5.

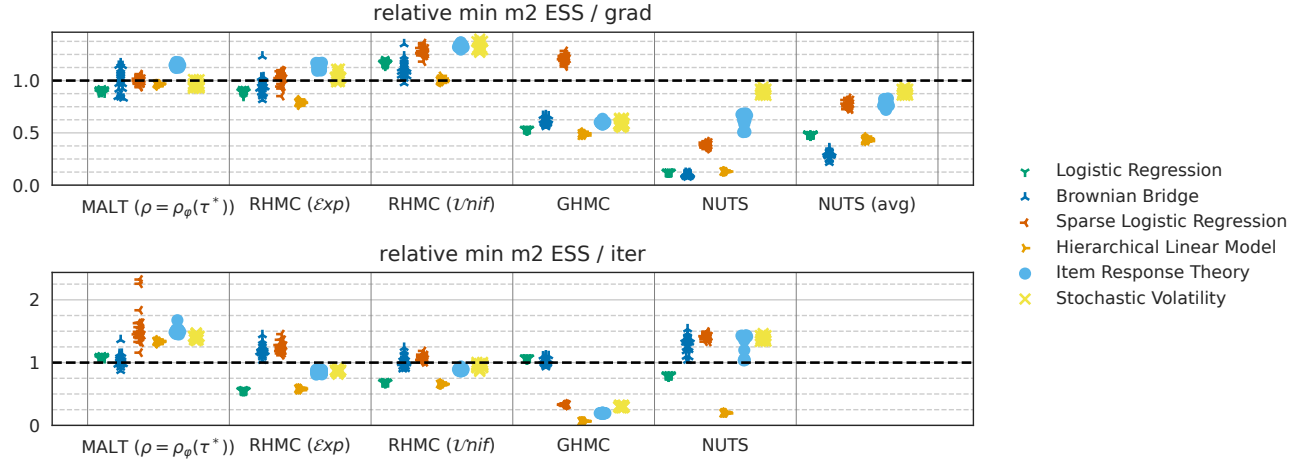


Figure 4: Normalized run values used to derive Figure 2. Each configuration shows one point per run. Horizontal jitter added for clarity.

Appendix C RAW ESS VALUES

In this section, we show the raw ESS/grad (Table 3) and ESS/iter (Table 4) numbers used to build Figure 2; see Section 5. We also show the raw run values in Figure 4. We use the shorthand notation $\rho^* \triangleq \rho_\varphi(\tau^*)$.

Table 3: Minimum ESS/grad for estimating the centered second moment across algorithms and models. Values are 10th percentile across 20 runs.

	MALT ($\rho = \rho^*$)	MALT ($\rho = 1$)	RHMC ($\mathcal{E}xp$)	RHMC ($Unif$)	GHMC	NUTS	NUTS (avg)
Logistic Regression	9.98e-2	1.10e-1	9.81e-2	1.30e-1	5.88e-2	1.36e-2	5.31e-2
Brownian Bridge	1.06e-2	1.15e-2	1.08e-2	1.30e-2	7.14e-3	1.02e-3	3.06e-3
Sparse Logistic Regression	5.34e-3	5.32e-3	5.20e-3	6.85e-3	6.45e-3	1.98e-3	4.04e-3
Hierarchical Linear Model	5.61e-3	5.76e-3	4.57e-3	5.81e-3	2.80e-3	7.50e-4	2.48e-3
Item Response Theory	8.97e-3	7.72e-3	8.81e-3	1.04e-2	4.66e-3	4.10e-3	5.86e-3
Stochastic Volatility	1.23e-2	1.28e-2	1.31e-2	1.69e-2	7.42e-3	1.13e-2	1.13e-2

Table 4: Minimum ESS/iter for estimating the centered second moment across algorithms and models. Values are 10th percentile across 20 runs.

	MALT ($\rho = \rho^*$)	MALT ($\rho = 1$)	RHMC ($\mathcal{E}xp$)	RHMC ($Unif$)	GHMC	NUTS
Logistic Regression	5.99e-1	4.78e-1	3.05e-1	3.77e-1	5.88e-1	4.36e-1
Brownian Bridge	7.20e-2	6.88e-2	8.33e-2	7.02e-2	7.14e-2	8.24e-2
Sparse Logistic Regression	2.69e-1	1.91e-1	2.33e-1	2.06e-1	6.45e-2	2.77e-1
Hierarchical Linear Model	5.61e-1	4.20e-1	2.43e-1	2.78e-1	2.80e-2	8.31e-2
Item Response Theory	3.59e-1	2.39e-1	2.02e-1	2.13e-1	4.66e-2	2.60e-1
Stochastic Volatility	3.56e-1	2.55e-1	2.21e-1	2.33e-1	7.42e-2	3.51e-1

Appendix D HYPERPARAMETERS

In this section we list the main hyperparameters used in experiments. For detailed implementation details, please refer to the source code available at https://github.com/tensorflow/probability/tree/main/discussion/adaptive_malt.

Table 5: Hyperparameters used in experiments.

Hyperparameter	Meaning	Value
K	Number of parallel chains	128
N_{clip}	Number of MALA steps	100
N_{adapt}	Number of adaptation steps	5000
α^*	Target acceptance rate	0.8 (0.7 NUTS)
η	ADAM learning rate for $\log(h)$ and $\log(\tau)$	0.05
a	Learning rate for online estimates except for w	8
a_w	Learning rate for w	3
β_1	ADAM hyperparameter for $\log(\tau)$	0
β_2	ADAM hyperparameter for $\log(\tau)$	0.95
ϵ	ADAM hyperparameter for $\log(\tau)$	10^{-8}
N_{ADVI}	Number of ADVI learning steps	2000
η_{ADVI}	ADVI learning learning rate	0.1