

---

# A Targeted Accuracy Diagnostic for Variational Approximations

---

Yu Wang  
Boston University

Mikołaj Kasprzak  
MIT and University of Luxembourg

Jonathan H. Huggins  
Boston University

## Abstract

Variational Inference (VI) is an attractive alternative to Markov Chain Monte Carlo (MCMC) due to its computational efficiency in the case of large datasets and/or complex models with high-dimensional parameters. However, evaluating the accuracy of variational approximations remains a challenge. Existing methods characterize the quality of the whole variational distribution, which is almost always poor in realistic applications, even if specific posterior functionals such as the component-wise means or variances are accurate. Hence, these diagnostics are of practical value only in limited circumstances. To address this issue, we propose the *T*Argeted *D*iagnostic for *D*istribution *A*pproximation *A*ccuracy (TADDAA), which uses many short parallel MCMC chains to obtain lower bounds on the error of each posterior functional of interest. We also develop a reliability check for TADDAA to determine when the lower bounds should not be trusted. Numerical experiments validate the practical utility and computational efficiency of our approach on a range of synthetic distributions and real-data examples, including sparse logistic regression and Bayesian neural network models.

## 1 INTRODUCTION

Bayesian inference is widely used due to its flexibility to handle arbitrary models, the desirable decision-theoretic properties of the provided point estimates, and its coherent uncertainty quantification (Robert et al., 2007). However, posterior point estimates and uncertainties usually cannot be calculated analytically, so in practice users of Bayesian methods rely on approximation algorithms. Variational inference (VI) has become an increasingly popular alternative

to Markov chain Monte Carlo (MCMC) due to its computational efficiency and the availability of black-box versions that can be applied to nearly arbitrary models (Jordan et al., 1999; Wainwright et al., 2008; Hoffman et al., 2013; Kingma and Welling, 2013; Blei et al., 2017; Kucukelbir et al., 2017; Agrawal et al., 2020).

The computational efficiency of variational inference arises from optimizing a divergence between the exact posterior and a distribution constrained to a tractable approximating family. The trade-off is that even the optimal distribution within this family may be a poor approximation to the posterior, so in general there is no guarantee that variational point estimates and uncertainties will be accurate. Therefore, it is important to develop diagnostic methods that quantify the error introduced by using a variational approximation. However, existing diagnostic methods either lack interpretability (Gorham and Mackey, 2017, 2015; Wada and Fujisaki, 2015; Chee and Toulis, 2018; Burda et al., 2015; Domke and Sheldon, 2018), or do not provide useful information when the parameter space is high-dimensional (Yao et al., 2018; Xing et al., 2020; Huggins et al., 2020). For example, Yao et al. (2018) propose using Pareto smoothed importance sampling (PSIS)  $\hat{k}$ , which essentially quantifies how effective the approximation is as a proposal distribution for importance sampling – not whether it will provide good approximations to specific quantities of interest. The kernel Stein discrepancy (KSD; Gorham and Mackey, 2017, 2015) can be used to evaluate the discrepancy between the approximate and the true distribution but its value is not readily interpretable and depends on a distant dissipativity condition that is not always easy to verify in practice. VI quality can also be accessed by monitoring the running average of changes in the estimated objective, which is usually the evidence lower bound (ELBO) (Wada and Fujisaki, 2015; Chee and Toulis, 2018; Burda et al., 2015; Domke and Sheldon, 2018). However, like the KSD, the ELBO lacks interpretability because it includes an unknown constant offset (the marginal likelihood), which is hard to estimate well unless the approximation is very accurate. Moreover, all these methods attempt to characterize the accuracy of the full variational distribution, which in practice is usually poor even if specific summaries such as posterior means, variances, or credible intervals are accurate. Therefore, in practice, they provide limited useful diagnostic information

in the types of complex models where variational methods offer the greatest computational savings.

In this paper, we propose a diagnostic that is applicable to high-dimensional parameter spaces and can provide lower bounds on the error of specific posterior summaries, including marginal means and covariances. Our approach is to run many short MCMC chains with stationary distribution equal to the posterior  $\pi$ . These chains are initialized with independent samples from a posterior approximation  $\hat{\pi}^{(0)}$ , which, after  $T$  iterations, results in an improved empirical distribution  $\hat{\pi}^{(T)}$  that is closer to  $\pi$ . If  $\hat{\pi}^{(0)}$  is a poor approximation to  $\pi$ , then the improved distribution  $\hat{\pi}^{(T)}$  obtained using MCMC can be significantly different from  $\hat{\pi}^{(0)}$ ; otherwise, the distribution will not shift. Thus, we can obtain lower bounds on the error of a posterior summary by computing confidence intervals for how much that summary differs under  $\hat{\pi}^{(0)}$  and  $\hat{\pi}^{(T)}$ . Our approach might fail to provide meaningful results, however, if the Markov kernel used mixes poorly. To ameliorate this issue, we jointly adapt the Markov chains at each iteration. We show that if the number of chains is large, they remain essentially independent despite the joint adaptation procedure; hence, our confidence intervals remain valid. We also propose a simple correlation-based check that can detect poorly mixing chains, and hence warn the user against trusting the diagnostic. Numerical experiments validate the practical utility and computational efficiency of our approach on a range of synthetic distributions and real-data examples.<sup>1</sup>

## 2 PRELIMINARIES

Modern Bayesian statistics relies heavily on models with difficult-to-compute posterior distributions, and hence requires general-purpose algorithms to approximate expectations with respect to those distributions. Throughout we abuse notation and use  $\pi$  to denote the posterior distribution and the corresponding density.

**Markov Chain Monte Carlo.** Markov chain Monte Carlo (MCMC) sampling methods provide a general-purpose framework for obtaining samples that are asymptotically exact (in runtime). To construct a Markov kernel with the desired stationary distribution, a standard approach is to correct an arbitrary proposal kernel using the Metropolis–Hastings (MH) accept-reject procedure (Metropolis et al., 1953). For  $x \in \mathbb{R}^d$ , let  $Q_\psi(x, dy)$  denote the proposal kernel parameterized by  $\psi$  with current state  $x$  and corresponding density  $q_\psi(x, y)$ . A proposed state  $Y \sim Q_\psi(x, \cdot)$  is accepted with probability

$$\alpha(x, Y) = \min \left\{ 1, \frac{\pi(Y)q_\psi(Y, x)}{\pi(x)q_\psi(x, Y)} \right\}.$$

<sup>1</sup>A Python implementation of TADDAA and code to reproduce all of our experiments is available at <https://github.com/TARPS-group/TADDAA>.

The Metropolis–Hastings Markov transition kernel is therefore given by

$$P_\psi(x, dy) := \alpha(x, y)Q_\psi(x, dy) + r_\psi(x)\delta_x(dy),$$

where  $r_\psi(x) := 1 - \int \alpha(x, y)Q_\psi(x, dy)$  is the rejection probability and  $\delta_x$  is a Dirac measure at  $x$ .

**Variational Inference.** Variational inference (VI) provides a potentially faster alternative to MCMC when models are complex and/or the dataset size is large (Blei et al., 2017). Variational inference aims to minimize some measure of discrepancy  $\mathcal{D}_\pi(\cdot)$  over a tractable family  $\mathcal{Q}$  of potential approximating distributions (Wainwright et al., 2008), resulting in a posterior approximation given by

$$\hat{\pi} = \arg \min_{\xi \in \mathcal{Q}} \mathcal{D}_\pi(\xi).$$

We will focus on the *Kullback–Leibler (KL) divergence*

$$\mathcal{D}_\pi(\xi) = \text{KL}(\xi | \pi) := \int \log \left( \frac{d\xi}{d\pi} \right) d\xi$$

as the discrepancy since it is the most widely used: the unknown marginal likelihood does not affect the optimization, computing gradients requires estimating expectations only with respect to distribution within a tractable family  $\mathcal{Q}$ , and optimization is relatively stable (Blei et al., 2017).

**Alternative Non-exact Approximation Methods.** While our experiments focus on variational inference, numerous other non-exact posterior approximation methods are widely used in practice such as expectation-propagation (Minka, 2001; Bishop, 2006), the Laplace approximation (Gelman et al., 1995; Bishop, 2006), and the iterated nested Laplace approximation (Rue et al., 2009; Gomez-Rubio, 2020). Our diagnostic applies equally well to all of these methods.

## 3 AN MCMC-BASED ACCURACY DIAGNOSTIC

Suppose we have an approximation  $\hat{\pi}^{(0)}$  to the posterior density  $\pi$ . Given a sample  $X^{(0)} \sim \hat{\pi}^{(0)}$ , we could improve it by applying a Markov kernel  $P_\psi(x, dy)$  with invariant distribution  $\pi$ , resulting in a sample  $X^{(1)} \sim P_\psi(X^{(0)}, \cdot)$  with a distribution  $\hat{\pi}^{(1)}$  that is typically closer to  $\pi$  than  $\hat{\pi}^{(0)}$  if  $\hat{\pi}^{(0)} \neq \pi$ . Continuing to run the Markov chain for  $T$  total steps results in a sample  $X^{(T)}$  with distribution  $\hat{\pi}^{(T)}$  that, when  $\hat{\pi}^{(0)}$  is a poor approximation to  $\pi$ , can be substantially closer to  $\pi$ ; see Fig. 1 for an illustration. More specifically, consider some posterior functional of interest such as a mean or variance, which we generically denote by  $\mathcal{F}$ . For example,  $\mathcal{F}(\pi) = \int x_i \pi(dx)$  would be the mean of the  $i$ th component. If the initial approximation error  $\varepsilon^{(0)}(\mathcal{F}) := \mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\pi)$  is large, then a well-chosen Markov kernel will result in the improved approximation error  $\varepsilon^{(T)}(\mathcal{F}) :=$




Figure 1: A cartoon plot for MCMC-based accuracy diagnostic. Even if the Markov chains do not reach the stationary distribution, they can provide a lower bound on error in the posterior approximation.

$\mathcal{F}(\hat{\pi}^{(T)}) - \mathcal{F}(\pi)$  being substantially smaller. In this case,  $|\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\hat{\pi}^{(T)})|$  measures the decrease in the initial approximation error, and hence provides a lower bound on  $|\varepsilon^{(0)}(\mathcal{F})|$ . If, on the other hand, the initial approximation error is small, then  $|\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\hat{\pi}^{(T)})|$  should remain small as well.

While in principle the just-described approach is appealing, there are two immediate problems. First, we do not have direct access to  $\hat{\pi}^{(T)}$ . Second,  $P_\psi(x, dy)$  may not be particularly efficient unless the proposal parameter  $\psi$  is well-chosen – which is nearly impossible to do *a priori*. To address both of these problems, we propose to run  $N$  Markov chains in parallel and, following standard MCMC practice, adapt the proposal parameter at each iteration (Andrieu and Thoms, 2008). Denote the state of the  $j$ th chain after iteration  $t$  by  $X_j^{(t)}$  and the proposal parameter used at iteration  $t$  by  $\psi^{(t-1)}$ . We write  $X_{1:N}^{(t)} = (X_j^{(t)})_{j=1}^N$  and similarly for  $X_j^{(1:t)}$ ,  $X_{1:N}^{(1:t)}$ , and other quantities. In the case of a Metropolis–Hastings kernel, at iteration  $t$ , for each  $j \in \{1, \dots, N\}$ , we generate proposals  $Y_j^{(t+1)} \sim Q_{\psi^{(t)}}(X_j^{(t)}, \cdot)$ , then accept or reject the proposal to obtain  $X_j^{(t+1)}$ . Next, we can use *inter-chain adaptation* (INCA; Craiu et al., 2009), which uses all the samples available up to the current time to update the proposal parameter to  $\psi^{(t+1)} = H(\psi^{(1:t)}, X_{1:N}^{(1:t+1)}, Y_{1:N}^{(1:t+1)})$ . We call  $H$  the *adaptation function*.

Finally, we can use the  $N$  samples from iteration  $T$  to construct a  $(1 - \alpha)$ -confidence interval  $(\ell_{\mathcal{F}}, u_{\mathcal{F}})$  for  $\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\hat{\pi}^{(T)})$ , and hence a high-confidence lower bound on  $|\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\pi)|$ , if the following condition holds:

**Condition 1.** *The map  $t \mapsto |\varepsilon^{(t)}(\mathcal{F})|$  is non-increasing and the sign of  $\varepsilon^{(t)}(\mathcal{F})$  does not depend on  $t$ .*

### Algorithm 1 TADDAA

**Input:** log density of the target  $\pi$ ,  
 approximating distribution  $\hat{\pi}^{(0)}$ ,  
 functional of interest  $\mathcal{F}$ ,  
 proposal kernel  $Q_\psi(x, dy)$  [Section 3.1],  
 initial proposal parameter  $\psi^{(0)}$  [Section 3.1],  
 number of Markov chains  $N$  [Eq. (6)],  
 number of iterations  $T$  [Eq. (7) or (8)]

- 1:  $X_j^{(0)} \sim \hat{\pi}^{(0)}$  for  $j = 1, \dots, N$
- 2: **for**  $t = 0$  **to**  $T - 1$  **do**
- 3:     **for**  $j = 1$  **to**  $N$  **do**
- 4:          $Y_j^{(t+1)} \sim Q_{\psi^{(t)}}(X_j^{(t)}, \cdot)$
- 5:          $\alpha_j^{(t)} = \alpha(X_j^{(t)}, Y_j^{(t)})$
- 6:          $X_j^{(t+1)} = \begin{cases} Y_j^{(t)}, & \text{with probability } \alpha_j^{(t)} \\ X_j^{(t)}, & \text{with probability } 1 - \alpha_j^{(t)} \end{cases}$
- 7:     **end for**
- 8:      $\psi^{(t+1)} = H(\psi^{(1:t)}, X_{1:N}^{(1:t+1)}, Y_{1:N}^{(1:t+1)})$  [Eq. (5)]
- 9: **end for**
- 10: Use  $\hat{\pi}^{(0)}$  (or  $X_{1:N}^{(1)}$ ) and  $X_{1:N}^{(T)}$  to compute a  $(1 - \alpha)$ -confidence interval  $(\ell_{\mathcal{F}}, u_{\mathcal{F}})$  for  $\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\hat{\pi}^{(T)})$
- 11: Compute correlation check  $\rho_{\max}^2(T)$  to verify reliability of the diagnostic [Section 3.4]
- 12: **return**  $B_{\mathcal{F}}$  as defined in Eq. (1) and  $\rho_{\max}^2(T)$

Under Condition 1, with probability  $1 - \alpha$ ,

$$\begin{aligned} \varepsilon^{(0)}(\mathcal{F}) &\geq |\mathcal{F}(\hat{\pi}^{(0)}) - \mathcal{F}(\hat{\pi}^{(T)})| \\ &\geq \begin{cases} 0 & \text{if } \ell_{\mathcal{F}} \leq 0 \leq u_{\mathcal{F}} \\ \min(|\ell_{\mathcal{F}}|, |u_{\mathcal{F}}|) & \text{otherwise} \end{cases} \quad (1) \\ &= \mathbb{1}\{0 \notin (\ell_{\mathcal{F}}, u_{\mathcal{F}})\} \times \min(|\ell_{\mathcal{F}}|, |u_{\mathcal{F}}|) \\ &=: B_{\mathcal{F}}. \end{aligned}$$

Algorithm 1 summarizes the key steps of our proposed *Targeted Diagnostic for Distribution Approximation Accuracy* (TADDAA) for computing the error lower bound  $B_{\mathcal{F}}$ . References to specific implementation recommendations (all described and justified in detail below) are in magenta.

**Remark 3.1.** *While in general we expect Condition 1 to hold, it is possible that it could be violated for target distributions with complex geometries. Although we only consider reversible Markov kernels, it is plausible that non-reversible Markov kernels could also lead to violations.*  $\triangleleft$

We will construct confidence intervals under the following condition.

**Condition 2.** *The elements of  $X_{1:N}^{(T)}$  are identically distributed, have finite second moments, and are pairwise independent.*

We will focus on diagnosing the accuracy of means, variances, and quantiles. But TADDAA can be straightforwardly applied to other functionals of interest. For

$t \in \{0, \dots, T\}$ , define the mean vector  $\mu^{(t)} = (\mu_1^{(t)}, \dots, \mu_d^{(t)})^\top := \mathbb{E}_{X \sim \hat{\pi}^{(t)}}(X)$ , and marginal standard deviations  $\sigma_i^{(t)} := \text{stdev}_{X \sim \hat{\pi}^{(t)}}(X_i)$  for  $i = 1, \dots, d$ .

**Marginal mean.** For the mean functional  $\mathcal{F}_{\text{mean}}^i(\pi) := \mu_i(\pi) := \int x_i \pi(dx)$ , under Condition 2, a  $(1 - \alpha)$  confidence interval for  $\mu_i^{(T)} - \mu_i^{(0)}$  is given by (Mendenhall et al., 2012)

$$\bar{x}_i^{(T)} - \mu_i^{(0)} \pm \frac{s_i^{(T)}}{\sqrt{N}} t_{N-1}(1 - \alpha/2), \quad (2)$$

where  $\bar{x}_i^{(T)} = \frac{1}{N} \sum_{j=1}^N x_{j,i}^{(T)}$  is the sample mean,  $s_i^{(T)} = \frac{1}{N-1} \left\{ \sum_{j=1}^N (x_{j,i}^{(T)} - \bar{x}_i^{(T)})^2 \right\}^{1/2}$  is the sample standard deviation, and  $t_{N-1}(q)$  is the  $q$  quantile of a  $t$  distribution with  $N - 1$  degrees of freedom.

**Marginal variance.** For the variance functional  $\mathcal{F}_{\text{var}}^i(\pi) := \sigma_i^2 := \text{Var}_{X \sim \pi}(X_i)$ , under Condition 2, a  $(1 - \alpha)$  confidence interval for  $2 \log(\sigma_i^{(T)}/\sigma_i^{(0)})$  is given by (Mendenhall et al., 2012)

$$\left[ \log \left( \frac{(N-1)s_i^{(T)}}{(\sigma_i^{(0)})^2 \chi_{N-1}^2(1-\alpha/2)} \right), \log \left( \frac{(N-1)s_i^{(T)}}{(\sigma_i^{(0)})^2 \chi_{N-1}^2(\alpha/2)} \right) \right], \quad (3)$$

where  $\chi_{N-1}^2(q)$  is the  $q$  quantile of a  $\chi^2$  distribution with  $N - 1$  degrees of freedom.

**Marginal quantile.** For the  $p$ -quantile functional

$$\mathcal{F}_p^i(\pi) = Q_p^i(\pi) := \inf\{q \in \mathbb{R} : P_{X \sim \pi}(X_i \leq q) \geq p\}, \quad (4)$$

under Condition 2, a  $(1 - \alpha)$  confidence interval for  $Q_p^i(\hat{\pi}^{(T)}) - Q_p^i(\hat{\pi}^{(0)})$  is given by (Hahn and Meeker, 2011)

$$\left[ X_{(l),i}^{(T)} - Q_p^i(\hat{\pi}^{(0)}), X_{(u),i}^{(T)} - Q_p^i(\hat{\pi}^{(0)}) \right],$$

where  $X_{(1),i}^{(T)} \leq X_{(2),i}^{(T)} \leq \dots \leq X_{(N),i}^{(T)}$  are the order statistics of the samples  $X_{1:N}^{(T)}$ ,  $l = B(\alpha/2, N, p)$ ,  $u = B(1 - \alpha/2, N, p) + 1$ , and  $B(q, N, p)$  is the  $q$  quantile of a binomial distribution with parameters  $N$  and  $p$ .

While we have now described the key components of TAD-DAA, a practical and reliable implementation requires addressing a number of issues:

1. **Proposal kernel and adaptation scheme.** The choice of MCMC proposal kernel and adaptation function  $H$  can have a dramatic effect on sampling efficiency and thus the accuracy of TAD-DAA. We propose to use gradient-based proposals that scale well to high-dimensional problems and focus on adapting the step size parameter of these algorithms. (Section 3.1)

2. **Inter-chain dependence.** INCA introduces dependence between the Markov chains, which could invalidate the confidence intervals. However, we show that the samples from the Markov chains at any fixed iteration satisfy a *propagation-of-chaos* property, which implies that the chains are asymptotically independent (in the limit  $N \rightarrow \infty$ ), so an asymptotic version of Condition 2 still holds. (Section 3.2)

3. **Choice of  $T$  and  $N$ .** The length of each Markov chain  $T$  and the number of Markov chains  $N$  must be chosen such that the diagnostic is computationally efficient while still (i) having  $T$  large enough that the improved approximation  $\hat{\pi}^{(T)}$  becomes sufficiently different from the initial approximation  $\hat{\pi}^{(0)}$  and (ii) having  $N$  large enough that the confidence intervals are sufficiently small to provide meaningful diagnostic information. (Section 3.3)

4. **Checking the reliability of the diagnostic.** Since it is possible the proposal kernel will be inefficient despite adaptation, we propose a correlation-based reliability check for our diagnostic. (Section 3.4)

### 3.1 Markov Kernels and Adaptation

We consider four possible Markov kernels: random walk Metropolis–Hastings (RWMH), Metropolis-adjusted Langevin algorithm (MALA), the Barker algorithm (Barker), and Hamiltonian Monte Carlo (HMC), which we describe in detail in Appendix A. All four kernels rely on a step-size parameter  $h$ . The latter three algorithms exploit gradient information about  $\log \pi$ , and thus have superior efficiency when  $d$  is large. High-dimensional sampling efficiency can be formalized using the theory of *optimal scaling* (Roberts and Rosenthal, 2001). To guarantee sampling efficiency and obtain well-behaved limiting behavior, optimal scaling results show that step size  $h$  needs to scale as  $\Theta(d^{-\gamma})$ , which results in a mixing time (i.e., number of iterations to obtain an effectively independent sample) of  $\Theta(d^\gamma)$ . Hence, the smaller  $\gamma$  is the more efficient the kernel. The values of  $\gamma$  for the kernels we consider are 1 for RWMH (Gelman et al., 1997),  $1/3$  for MALA (Roberts and Rosenthal, 1998),  $1/3$  for Barker (Livingstone and Zanella, 2022), and  $1/4$  for HMC (Beskos et al., 2013). Optimal scaling theory also provides the optimal asymptotic acceptance rates ( $\bar{\alpha}_*$ ) for each kernel, which are 0.234 for RWMH, 0.4 for Barker, 0.574 for MALA, and 0.651 for HMC. While HMC has the best scaling with  $\gamma = 1/4$ , the efficiency of Barker is more robust to the precise step size (and hence precise acceptance rate) (Livingstone and Zanella, 2022), which makes it attractive for our use-case where it may not be possible to obtain the optimal step size since adaptation occurs over a relatively small number of iterations  $T$ .

We adapt the step size so that the acceptance rate approaches the optimal asymptotic acceptance rate  $\bar{\alpha}_*$  (Andrieu and

Thoms, 2008). Let  $\bar{\alpha}^{(t)} := N^{-1} \sum_{j=1}^N \alpha_j^{(t)}$  denote the average acceptance probability at iteration  $t$ . Letting  $\psi^{(t)} = \log h^{(t)}$  be the log of the step size, the adaptation function is given by

$$\begin{aligned} H(\psi^{(1:t)}, X_{1:N}^{(1:t+1)}, Y_{1:N}^{(1:t+1)}) \\ = \psi^{(t)} + \frac{1}{\sqrt{t+1}} (\bar{\alpha}^{(t)} - \bar{\alpha}_*). \end{aligned} \quad (5)$$

Compared with running many parallel chains independent of each other, using a joint adaptation scheme improves adaptation speed and ensures the samples at each iteration follow the same distribution (Rosenthal, 2000; Solonen et al., 2012). The sequence of proposal parameters  $(\psi^{(t)})_{t=0}^{\infty}$  converge under the standard Robbins–Monro conditions (see Andrieu and Thoms (2008, Section 4) for a related discussion). In order to prove convergence to the stationary distribution of the resulting Markov chain, one needs to verify an additional assumption, such as uniform ergodicity Roberts and Rosenthal (2007, Theorem 1) or one of the weaker sufficient conditions proposed, for instance in Atchadé and Rosenthal (2005, Assumption 3.1) or Green et al. (2015, Section 2.4).

Following Roberts and Rosenthal (1998), we initialize the step sizes to  $h^{(0)} = 2.4^2/d$  for RWMH, and  $h^{(0)} = 2.4^2/d^{1/3}$  for MALA and Barker, and  $h^{(0)} = 2.4^2/d^{1/4}$  for HMC. All of the Markov kernels we consider can make use of a pre-conditioning matrix, which is beneficial when components have different scales (Roy and Zhang, 2022; Stramer and Roberts, 2007). While it is possible to adapt the pre-conditioning matrix in addition to the step size in practice we found it sufficed to keep the pre-conditioning matrix fixed and equal to the covariance of  $\hat{\pi}^{(0)}$ . If the covariance is not available in closed form we use the sample covariance of the samples  $X_{1:N}^{(0)}$ .

### 3.2 Asymptotic Independence of Adapted Markov Chains

Due to the use of adaptation, the Markov chains  $X_{1:N}^{(t)}$  are not independent once  $t > 1$ , so the final samples  $X_{1:N}^{(T)}$  violate the independence requirement of Condition 2. However, we can instead ensure that the Markov chains are *asymptotically* independent in the following sense:

**Definition 3.2.** Let  $X_{N,1:N} = (X_{N,1}, \dots, X_{N,N})$  denote a random vector. The sequence of random vectors  $\{X_{N,1:N}\}_{N=1}^{\infty}$  is  $\bar{\nu}$ -chaotic if, for any  $r \in \mathbb{N}$  and any bounded continuous real-valued functions  $g_1, g_2, \dots, g_r$ ,

$$\lim_{N \rightarrow \infty} \mathbb{E}_{X_{N,1:N}} \left\{ \prod_{i=1}^r g_i(X_{N,i}) \right\} = \prod_{i=1}^r \int g_i(x) \bar{\nu}(dx).$$

In particular, we show that the Markov chains are chaotic after any fixed number of iterations.

**Theorem 3.3.** Under Assumption 3.4 (given below), for any  $t \in \mathbb{N}$ , there exists a probability distribution  $\bar{\nu}^{(t)}$  such that the sequence  $\{X_{1:N}^{(t)}\}_{N=1}^{\infty}$  is  $\bar{\nu}^{(t)}$ -chaotic.

The proof can be found in Appendix B. Crucially, the asymptotics we require for Theorem 3.3 are in the number of Markov chains  $N$ , not the number of iterations  $T$ . Thus, for any fixed  $T$ , Condition 2 holds asymptotically in  $N$ .

To write our assumptions, note that if we treat the proposed states  $Y_{1:N}^{(t)}$  as additional random variables, we can write the MCMC update as

$$(X_j^{(t+1)}, Y_j^{(t+1)}) \sim T(X_j^{(t)}, Y_j^{(t)}, h_N^{(t)}, \cdot, \cdot)$$

for an appropriate Markov kernel  $T(x, y, h, dx', dy')$ , where  $h$  is the step size.

**Assumption 3.4.** (a) The proposal probability density  $q_h(y, x)$  is continuous with respect to  $(x, y, h)$ .

(b) The target distribution has a continuous probability density function  $\pi(\cdot)$ .

(c) Samples generated from the Markov transition kernel  $T(x, y, h, \cdot, \cdot)$  satisfy  $\mathbb{E}\|X_j^{(t)}\|^2 < \infty$  and  $\mathbb{E}\|Y_j^{(t)}\|^2 < \infty$  for any  $t \in \mathbb{N}$ .

These conditions are all quite mild. All of the proposals used in our experiments satisfy Assumption 3.4.(a). Assumption 3.4.(b) is required to use all three gradient-based proposals (MALA, Barker, and HMC). Assumption 3.4.(c) just requires that for any fixed time  $t \in \mathbb{N}$ , the generated samples have finite second moment.

**Remark 3.5.** Theorem 3.3 could be generalized to targets with discrete components if the proposal distribution is absolutely continuous with respect to an appropriate reference measure and the topology of the space on which the target is defined respects the discrete and continuous components of the target. (For example, if the target distribution is defined on  $\mathbb{R}$  and has an atom at zero, then  $\{0\}$  should be an open set.) In that case, Definition 3.2 would need to be adapted in order to account for the change in the underlying topology, and so will the details of the proof of Theorem 3.3.  $\triangleleft$

### 3.3 Length and Number of Markov Chains

Having established the validity of Algorithm 1, we now turn to the question of how to ensure the computational cost is not prohibitive and the diagnostic is accurate. Given a target distribution and Markov kernel, the computational cost is determined by the number of Markov chains  $N$  and iterations  $T$ . We address each in turn.

**Number of Markov Chains.** The number of Markov chains must be sufficiently large that the confidence intervals are small enough to detect meaningful errors. Our

approach is to choose  $N$  to be the smallest value that satisfies the user’s tolerance level for the margin of error. In the case of mean estimation, it follows from Eq. (2) that the error for estimating  $\mu_i^{(T)}$  is given by  $t_{N-1}(\alpha/2) s_i^{(T)} / \sqrt{N}$ . Similarly, the log variance error can also be determined by Eq. (3). For the mean we use the relative error normalized by the standard deviation, as this accounts for the relevant scale of the error. Hence, if the user’s tolerance is  $\delta_{\text{mean}}$  for relative mean error and  $\delta_{\text{var}}$  for log variance error<sup>2</sup>, then the required sample size (number of Markov chains) is

$$N = \max(N_{\text{mean}}, N_{\text{variance}}), \quad (6)$$

where

$$N_{\text{mean}} := \min \left\{ n \in \mathbb{N} : \frac{t_{n-1}(\alpha/2)}{\sqrt{n}} \leq \delta_{\text{mean}} \right\},$$

$$N_{\text{variance}} := \min \left\{ n \in \mathbb{N} : \log \left( \frac{\chi_{n-1}^2(1-\alpha/2)}{\chi_{n-1}^2(\alpha/2)} \right) \leq \delta_{\text{var}} \right\}.$$

**Number of Iterations.** The number of iterations has to be large enough that the improved distribution  $\hat{\pi}^{(T)}$  is sufficiently different from a poor approximating distribution  $\hat{\pi}^{(0)}$ . As discussed earlier, the theory of optimal scaling shows that the Markov chain requires  $\Theta(d^7)$  iterations to mix. Thus, we propose to scale  $T$  with respect to dimension  $d$  in the same manner. For RWMH<sup>3</sup>, MALA, and Barker, we take

$$T = \lfloor cd^{1/3} \rfloor, \quad (7)$$

while for HMC we take

$$T = \lfloor cd^{1/4}/L \rfloor, \quad (8)$$

where  $L$  is the number of leapfrog steps in HMC. Scaling by  $L$  is necessary for HMC to ensure  $T$  accounts for the  $L$  gradient evaluations required per iteration. All other methods require 0 or 1 gradient evaluation per iteration. Based on numerical results we suggest  $c = 50$ . In Appendix C.1, we also run ablation studies on how the diagnostic result (several types of lower bound) would change with  $T$ , which shows that Eq. (7) and Eq. (8) are reasonable.

The results of Bhatia et al. (2022) suggest our choice of  $T$  will ensure the cost of TADDAA will not be too large compared to the cost of using black-box variational inference algorithms. Specifically, Bhatia et al. (2022, Theorem 1) shows that the number of gradient evaluations of stochastic optimization needed to obtain a constant optimization error scales as  $d^{1/3}$ . Hence, the computational cost (as measured by number of gradient evaluations) is  $\Theta(d^{1/3})$  for VI,  $\Theta(d^{1/3})$  for MALA and Barker, and  $\Theta(d^{1/4})$  for HMC. We verify the cost of TADDAA is not prohibitive in practice in our numerical experiments.

<sup>2</sup>The width of the quantile confidence interval is sample dependent, so we do not set a tolerance level for the quantile error.

<sup>3</sup>Based on optimal scaling behavior of RWMH, Eq. (7) may not be large enough to guarantee significant change in the distribution. So, we do not recommend RWMH.

**Remark 3.6.** TADDAA is essentially running MCMC only for a burn-in period. Crucially, the diagnostic remains valid and useful even if the chains remain far from stationarity as long as  $\hat{\pi}^{(T)}$  is closer to stationarity (i.e., it improves the approximating distribution  $\hat{\pi}^{(0)}$  somewhat). In challenging scenarios where moving closer to stationarity is hard even with best practices (e.g., INCA and preconditioning), the TADDAA error lower bounds can be close to zero, leading to uninformative diagnostic results. See Section 4.3 for a concrete example. We note also that there is a large literature studying the convergence rates of MCMC algorithms (Brooks et al., 2011; Roberts and Rosenthal, 2004).  $\triangleleft$

### 3.4 A Reliability Check for the Diagnostic

The reliability of TADDAA depends on the mixing behavior of the Markov chains:

- If the Markov chains are mixing well, we expect the diagnostic results can be trusted.
- If the Markov chains are not mixing well, diagnosis of a poor approximation can still be trusted but diagnosis of a good approximation may not be reliable. In the latter case, we should consider increasing the length of Markov chains or otherwise improving the Markov kernel.

The usual approach to checking MCMC mixing is some version of the  $\hat{R}$  diagnostic (Gelman and Rubin, 1992; Vats and Knudson, 2021; Vehtari et al., 2021). However, using  $\hat{R}$  requires a sufficiently large effective sample size for each Markov chain, which is not satisfied for the short Markov chains used in TADDAA. Instead, we propose to leverage the fact that we have many nearly independent chains to compute the correlation between the initial and final values of each chain. If the chains are mixing, then this correlation will be close to zero. Specifically we propose to use the worst-case correlation coefficient  $\rho_{\text{max}}^2(T) := \max_i \rho_i^2(T)$ , where

$$\rho_i^2(T) := \frac{\sum_{j=1}^N (X_{j,i}^{(0)} - \hat{\mu}_i^{(0)})(X_{j,i}^{(T)} - \hat{\mu}_i^{(T)})}{\sqrt{\sum_{j=1}^N (X_{j,i}^{(0)} - \hat{\mu}_i^{(0)})^2} \sqrt{\sum_{j=1}^N (X_{j,i}^{(T)} - \hat{\mu}_i^{(T)})^2}}.$$

If  $\rho_{\text{max}}^2(T)$  is close to 0 (we use a cutoff of 0.1), then the check passes; otherwise, the check fails. In addition to examining the value of  $\rho_{\text{max}}^2(T)$ , we can also examine plots of  $t$  versus  $\rho_{\text{max}}^2(t)$  and  $t$  versus  $\rho_i^2(t)$  to gain insight into, how quickly mixing is occurring, respectively, overall and for specific coordinates.

## 4 EXPERIMENTS

We now validate the efficacy and computational efficiency of TADDAA and the correlation-based reliability check. In our experiments, the number of Markov chains is  $N = 386$ ,

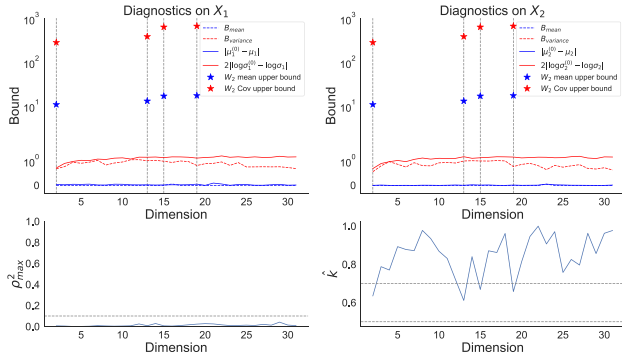


Figure 2: Diagnostics for mean-field VI approximation to correlated Gaussian targets, where TADDAA uses the Barker proposal. Here  $\mu_i$  and  $\sigma_i^2$  denote, respectively, the mean and variance of  $X_i$ .




Figure 3: Reliability checks for TADDAA using Barker applied to the mean-field VI approximation to correlated Gaussian targets.

which corresponds to the tolerance level  $\delta_{\text{mean}} = 0.1$  and  $\delta_{\text{var}} = 0.15$ . While we often present results for all four kernels discussed in Section 3.1 (RWMH, MALA, Barker, and HMC), we generally recommend using Barker because it has the efficiency of gradient-based kernels but is less sensitive to parameter tuning than MALA and HMC. For variational inference, we use the default settings of the bbvi function in the VIABEL package (Welandawe et al., 2022), which used a mean-field Gaussian approximation family. We set the maximum number of iterations to 200,000 but the optimization can terminate early if convergence is reached. We compare TADDAA to the  $\hat{k}$  diagnostic proposed in Yao et al. (2018) and the Wasserstein-based mean and variance upper bounds proposed in Huggins et al. (2020).

#### 4.1 Gaussian Model With Correlated Coordinates

We first investigate the performance of TADDAA in a high-dimensional setting where the variational mean-field approximation has correct mean but incorrect variance estimates, which is a common situation in practice. Specifically, we use a correlated Gaussian target  $\pi = \mathcal{N}(\mu, \Sigma)$ , where  $\mu \in \mathbb{R}^d$ ,  $\Sigma_{ii} = \sigma_i^2$ , and  $\Sigma_{ij} = \rho\sigma_i\sigma_j$  for  $i \neq j$ . We introduce correlation by setting  $\rho = 0.7$  and variance heterogeneity by setting  $\sigma_1^2 = 10$  and  $\sigma_i^2 = 1$  for  $i = 2, 3, \dots, d$ .

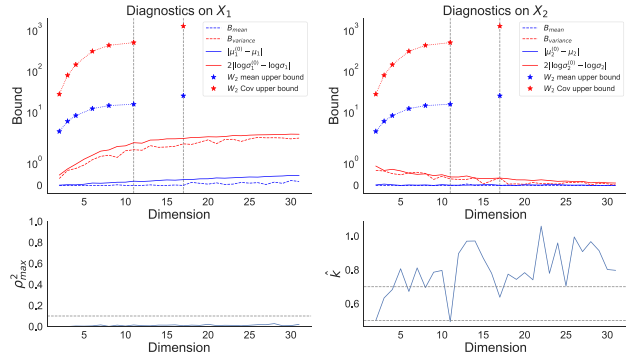


Figure 4: Diagnostics for Neal-funnel shape model, where TADDAA uses the Barker proposal. Here  $\mu_i$  and  $\sigma_i^2$  denote, respectively, the mean and variance of  $X_i$ .

Fig. 2 shows that the lower bounds  $B_{\text{mean}}$  and  $B_{\text{variance}}$  remain valid and useful even when  $d$  and  $k$  are large: TADDAA correctly captures that the mean estimates are accurate but the variance estimates are inaccurate. From Fig. 3, we can see that the Barker chains mix well in all cases. Figure 2 also provides a comparison to  $\hat{k}$  and the Wasserstein-based upper bounds proposed in Huggins et al. (2020, Theorem 3.4). The  $\hat{k}$  result suggests that the VI approximation is not reliable for almost any dimension (since  $\hat{k} > 0.7$ ). But if the functional of interest is the mean, VI should be considered reliable. Similarly, the Wasserstein upper bounds are too conservative by a factor of 10–1,000, incorrectly suggesting the VI approximation is extremely inaccurate. Moreover, the upper bounds are unavailable once  $\hat{k} > 0.7$  (Huggins et al., 2020). Overall, TADDAA provides much more precise and actionable information about the approximation quality than the other diagnostics.

#### 4.2 Neal-Funnel Shape Model

Next, we consider a more challenging target with a more complex geometry: Neal’s funnel distribution (Neal, 2003), which is similar to the geometries encountered in hierarchical models. For  $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ , let  $\sigma := \exp(x_1)$ . The funnel distribution is given by  $\pi(x) = \mathcal{N}(x_1 | 0, 1) \prod_{i=2}^{d-1} \mathcal{N}(x_i | 0, \sigma)$ . Hence, if  $X \sim \pi$ , then  $\text{Var}(X_i) = \exp(1/2)$  for  $i \in \{2, \dots, d\}$ . Since  $X_2, \dots, X_d$  are identically distributed, we focus on the accuracy of the mean-field VI approximations to the distributions of  $X_1$  and  $X_2$ .

Figure 4 shows that the lower bounds constructed using Barker are quite precise and do not degrade as the dimension increases. (See Appendix C.2 for results for quantiles.) Moreover the reliability check passes for all  $d$  (Fig. 13 in Appendix C.2). The  $\hat{k}$  diagnostic, however, is noisy and provides little insight. The Wasserstein-based mean and covariance upper bounds are orders-of-magnitude too large and unavailable for most  $d > 10$  because  $\hat{k} > 0.7$ .

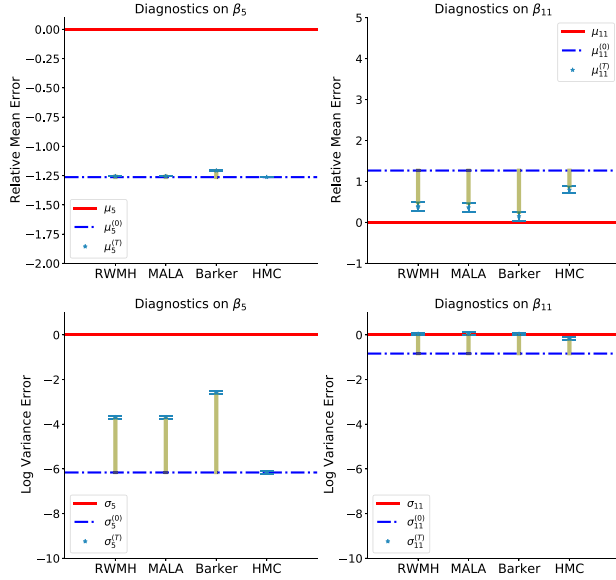


Figure 5: Diagnostics for logistic regression model on candy power ranking dataset for component  $\beta_5$  and  $\beta_{11}$ . Here  $\mu_i$  and  $\sigma_i$  denote, respectively, the mean and variance of  $\beta_i$ . To improve interpretability, we show the relative mean error  $(\mu_i^{(t)} - \mu_i)/\sigma_i$  and log variance error  $2 \log_{10}(\sigma_i^{(t)}) - 2 \log_{10}(\sigma_i)$ . (Recall that  $\mu_i$  and  $\sigma_i$  are the  $i$ th component mean and standard deviation of the target  $\pi$ .) The lengths of the gold lines give the corresponding relative mean and log variance error lower bounds.

### 4.3 Logistic Regression Model on Candy Power Ranking Dataset

Next, we validate TADDAA on a real-world dataset<sup>4</sup>: each observation  $(z_n, y_n)$  represents a different kind of candy, where  $z_n \in \mathbb{R}^{11}$  are the features, and  $y_n = 1$  if the candy is chocolate and 0 otherwise. We use a logistic regression model with  $y_n \sim \text{Bern}(\text{logit}^{-1}(\beta^\top z_n))$  and  $\beta_i \sim \mathcal{N}(0, I)$ .

Focusing on Barker, Fig. 5 shows that for  $\beta_5$ , the lower bound of the mean error is small but non-zero due to the fact that the variance is dramatically underestimated, leading to inefficient preconditioning (adaptive preconditioning did not significantly improve the lower bound). On the other hand, the variance error lower bound, while not tight, shows that the variance estimate is orders-of-magnitude too small. For  $\beta_{11}$ , where the variance is not so poorly estimated, both mean and variance lower bounds are very accurate. TADDAA uses about 2% as many gradient evaluations as the VI optimization ( $4.3 \times 10^4$  versus  $2 \times 10^6$ ), validating the computational efficiency of our approach. Figure 6 demonstrates the usefulness of the reliability checks, which correctly show that HMC chains fail to mix (due to a smaller number of




Figure 6: Reliability checks on TADDAA using different kernels applied to the mean-field VI approximation to logistic regression for candy classification.

iterations), while the chains for other kernels mix well.

### 4.4 Cancer Classification Using a Horseshoe Prior

We now consider a more challenging, higher-dimensional application to predicting leukemia using microarray data with  $n = 71$  observations and using 100 features chosen according to their  $\chi^2$  scores (Ray et al., 2020). We use a logistic regression model with a sparsity-inducing horseshoe prior (Piironen and Vehtari, 2017):

$$\begin{aligned} y \mid \beta &\sim \text{Bern}(\text{logit}^{-1}(X\beta)), \\ \beta_j \mid \tau, \lambda, c &\sim \mathcal{N}(0, \tau^2 \tilde{\lambda}_j^2), \\ \lambda_j &\sim C^+(0, 1), \quad \tau \sim C^+(0, \tau_0), \\ c^2 &\sim \text{InvGam}(2, 8), \end{aligned}$$

where  $y$  denotes the binary outcomes,  $X \in \mathbb{R}^{71 \times 100}$  is the features matrix,  $\tau > 0$  and  $\lambda > 0$  are global and local shrinkage parameters, and  $\tilde{\lambda}_j^2 = c^2 \lambda_j^2 / (c^2 + \tau^2 \lambda_j^2)$ . Hence, for our data the parameter dimensionality is  $d = 203$ . We would expect the VI approximation to be poor due to multimodality of the model. Focusing on two representative parameters, Fig. 7 shows that TADDAA using Barker correctly captures that the mean estimate for  $\lambda_{61}$  and variance estimate for  $\beta_0$  are inaccurate. On the other hand, for the accurate mean estimate for  $\beta_0$  and variance estimate for  $\lambda_{61}$ , the lower bounds are zero or close to zero. TADDAA uses about 28% as many gradient evaluations as the VI optimization ( $1.1 \times 10^5$  versus  $4 \times 10^5$ ). In addition, the reliability checks shown in Fig. 8 correctly reflect the worse performance of RWMH (due to  $d$  being larger) and HMC (due to a smaller number of iterations).

### 4.5 Bayesian Neural Network on Candy Power Ranking Dataset

Finally, we validate TADDAA on a Bayesian neural network model, which is difficult for mean-field VI to approximate accurately (Izmailov et al., 2021). We use the same dataset described in Section 4.3, but now it is fitted with a two-layer

<sup>4</sup><https://www.kaggle.com/datasets/fivethirtyeight/the-ultimate-halloween-candy-power-ranking>



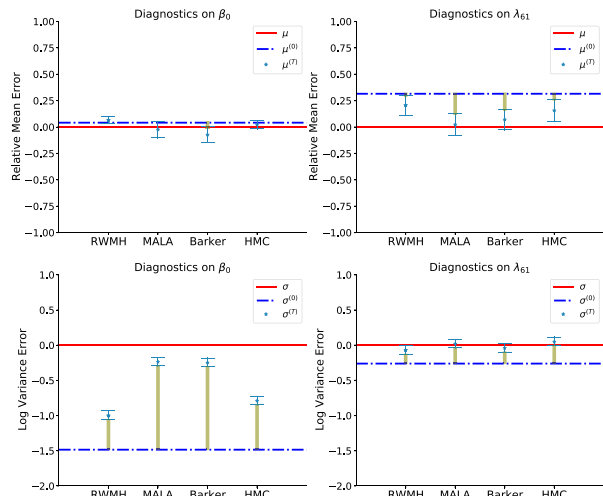


Figure 7: Diagnostics for  $\beta_0$  and  $\lambda_{61}$  for cancer classification using the horseshoe prior. See Fig. 5 for further explanations.




Figure 8: Reliability checks for TADDAA using different kernels applied to horseshoes prior applied to mean-field approximation for cancer classification.

Bayesian neural network with 5 units in each hidden layer:

$$y_n \sim \text{Bern}(\text{logit}^{-1}(\tanh(\tanh(\alpha^\top z_n) \cdot \beta) \cdot \gamma)),$$

where  $\alpha \in \mathbb{R}^{5 \times 11}$ ,  $\beta \in \mathbb{R}^{5 \times 5}$ ,  $\gamma \in \mathbb{R}^5$ ,  $\alpha_{i,j} \sim \mathcal{N}(0, 2)$ ,  $\beta_{i,j} \sim \mathcal{N}(0, 2)$  and  $\gamma_i \sim \mathcal{N}(0, 2)$ . Hence,  $x = (\alpha, \beta, \gamma)$  and  $d = 85$ . To quantify the classification quality for the Bayesian neural network, we can use the log loss (a.k.a., the cross-entropy loss)  $\mathcal{L}(x) = -\frac{1}{N} \sum_{n=1}^N \log p(y_n | x)$ . Similar to the definitions for the marginal mean and median, we can also use the mean and median of the log loss as the evaluation functional  $\mathcal{F}$ . Figure 9 shows that TADDAA correctly captures that the inaccuracy of the VI approximation: error lower bounds show that both the mean and median of the log loss are overestimated by mean-field VI. TADDAA uses about 12% as many gradient evaluations as the VI optimization ( $8.4 \times 10^4$  versus  $6.8 \times 10^5$ ). The reliability check in Fig. 15 confirms the superiority of Barker in high-dimensional problems.




Figure 9: Diagnostics of for Bayesian neural network model on candy power ranking dataset for component.  $\mu$  and  $Q(0.5)$  denote the mean and median for the cross-entropy loss.

## 5 CONCLUSION AND LIMITATIONS

Overall, our theory and experimental results show that TADDAA provides an efficient evaluation tool for variational inference, while also being applicable to other inexact posterior approximation methods like (integrated nested) Laplace approximations. TADDAA provides precise information about specific functionals of interest such as means and standard deviations, rather than a check on just the overall quality of an approximation – which in practice is often quite poor. However, there are a number of important limitations that users must keep in mind when employing TADDAA. While we have developed the correlation check to guard against a poor diagnostic, the check can fail due to, e.g., multimodality. On the other hand, such failures are not unique to our diagnostic and can affect, e.g., diagnostics for MCMC too. One approach to guard against this possibility is to run the VI optimization multiple times with diverse initializations, as is usually recommended for MCMC. Another limitation is that TADDAA is not appropriate in settings where it is only feasible to make a small number of passes over the entire dataset, so a method like stochastic variational inference (Hoffman et al., 2013) is used. Finally, it is possible that Condition 1 could be violated. Developing checks against this failure mode is an important direction for future work.

### Acknowledgments

M. Kasprzak was supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101024264. J. H. Huggins was partially supported by the National Institute of General Medical Sciences of the National Institutes of Health under grant number R01GM144963 as part of the Joint NSF/NIGMS Mathematical Biology Program. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

## References

- Agrawal, A., Sheldon, D. R., and Domke, J. (2020). Advances in black-box VI: Normalizing flows, importance weighting, and optimization. *Advances in Neural Information Processing Systems*, 33:17358–17369.
- Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373.
- Atchadé, Y. F. and Rosenthal, J. S. (2005). On adaptive Markov chain Monte Carlo algorithms. *Bernoulli*, 11(5):815 – 828.
- 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.
- Bhatia, K., Kuang, N. L., Ma, Y.-A., and Wang, Y. (2022). Statistical and Computational Trade-offs in Variational Inference: A Case Study in Inferential Model Selection. *arXiv preprint arXiv:2207.11208*.
- Billingsley, P. (2013). *Convergence of probability measures*. John Wiley & Sons.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*, volume 4. Springer.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.
- Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (2011). *Handbook of Markov chain Monte Carlo*. CRC press.
- Burda, Y., Grosse, R., and Salakhutdinov, R. (2015). Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*.
- Chee, J. and Toulis, P. (2018). Convergence diagnostics for stochastic gradient descent with constant learning rate. In *International Conference on Artificial Intelligence and Statistics*, pages 1476–1485. PMLR.
- Craiu, R. V., Rosenthal, J., and Yang, C. (2009). Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *Journal of the American Statistical Association*, 104(488):1454–1466.
- Domke, J. and Sheldon, D. R. (2018). Importance weighting and variational inference. *Advances in Neural Information Processing Systems*, 31.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.
- Gelman, A., Gilks, W. R., and Roberts, G. O. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120.
- Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, pages 457–472.
- Gomez-Rubio, V. (2020). *Bayesian inference with INLA*. CRC Press.
- Gorham, J. and Mackey, L. (2015). Measuring sample quality with Stein’s method. *Advances in Neural Information Processing Systems*, 28.
- Gorham, J. and Mackey, L. (2017). Measuring sample quality with kernels. In *International Conference on Machine Learning*, pages 1292–1301. PMLR.
- Green, P. J., Łatuszyński, K., Pereyra, M., and Robert, C. P. (2015). Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statistics and Computing*, 25(4):835–862.
- Hahn, G. J. and Meeker, W. Q. (2011). *Statistical intervals: a guide for practitioners*, volume 92. John Wiley & Sons.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. (2013). Stochastic Variational Inference. *Journal of Machine Learning Research*, 14:1303–1347.
- Huggins, J., Kasprzak, M., Campbell, T., and Broderick, T. (2020). Validated variational inference via practical posterior error bounds. In *International Conference on Artificial Intelligence and Statistics*, pages 1792–1802. PMLR.
- Izmailov, P., Vikram, S., Hoffman, M. D., and Wilson, A. G. G. (2021). What Are Bayesian Neural Network Posteriors Really Like? In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4629–4640. PMLR.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. M. (2017). Automatic differentiation variational inference. *Journal of Machine Learning Research*.
- Livingstone, S. and Zanella, G. (2022). The Barker proposal: Combining robustness and efficiency in gradient-based MCMC. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 84(2):496.
- Mendenhall, W., Beaver, R. J., and Beaver, B. M. (2012). *Introduction to probability and statistics*. Cengage Learning.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092.
- Minka, T. P. (2001). Expectation propagation for approximate Bayesian inference. In *Uncertainty in Artificial Intelligence*.

- Neal, R. (2011). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, pages 113–162.
- Neal, R. M. (2003). Slice sampling. *The Annals of Statistics*, 31(3):705–767.
- Piironen, J. and Vehtari, A. (2017). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2):5018–5051.
- Ray, S., Alshouli, K., Roy, A., AlGhamdi, A., and Agrawal, D. P. (2020). Chi-Squared Based Feature Selection for Stroke Prediction using AzureML. In *2020 Intermountain Engineering, Technology and Computing (IETC)*, pages 1–6. IEEE.
- Robert, C. P. et al. (2007). *The Bayesian choice: from decision-theoretic foundations to computational implementation*, volume 2. Springer.
- Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268.
- Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351 – 367.
- Roberts, G. O. and Rosenthal, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71.
- 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.
- Roberts, G. O. and Tweedie, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363.
- Rosenthal, J. S. (2000). Parallel computing and Monte Carlo algorithms. *Far East Journal of Theoretical Statistics*, 4(2):207–236.
- Roy, V. and Zhang, L. (2022). Convergence of position-dependent MALA with application to conditional simulation in GLMMs. *Journal of Computational and Graphical Statistics*, (just-accepted):1–31.
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):319 – 392.
- Sirignano, J. and Spiliopoulos, K. (2020). Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752.
- Solonen, A., Ollinaho, P., Laine, M., Haario, H., Tamminen, J., and Järvinen, H. (2012). Efficient MCMC for climate model parameter estimation: Parallel adaptive chains and early rejection. *Bayesian Analysis*, 7(3):715–736.
- Stramer, O. and Roberts, G. O. (2007). On Bayesian analysis of nonlinear continuous-time autoregression models. *Journal of Time Series Analysis*, 28(5):744–762.
- Sznitman, A.-S. (1991). *Topics in propagation of chaos*. Springer.
- Vats, D. and Knudson, C. (2021). Revisiting the gelman–rubin diagnostic. *Statistical Science*, 36(4):518–529.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., and Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: an improved  $\hat{R}$  for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2):667–718.
- Wada, T. and Fujisaki, Y. (2015). A stopping rule for stochastic approximation. *Automatica*, 60:1–6.
- Wainwright, M. J., Jordan, M. I., et al. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305.
- Welandawe, M., Andersen, M. R., Vehtari, A., and Huggins, J. H. (2022). Robust, Automated, and Accurate Black-box Variational Inference. *arXiv*, arXiv:2203.15945 [stat.ML].
- Xing, H., Nicholls, G., and Lee, J. K. (2020). Distortion estimates for approximate Bayesian inference. In *Conference on Uncertainty in Artificial Intelligence*, pages 1208–1217. PMLR.
- Yao, Y., Vehtari, A., Simpson, D., and Gelman, A. (2018). Yes, but did it work?: Evaluating variational inference. In *International Conference on Machine Learning*, pages 5581–5590. PMLR.

## Supplementary Materials for A Targeted Accuracy Diagnostic for Variational Approximations

---

### A METROPOLIS-HASTING KERNELS

We briefly summarize the kernels used in our experiments. Throughout this section let  $x \in \mathbb{R}^d$  denote the current state,  $h \in \mathbb{R}_+$  the step size, and  $G \in \mathbb{R}^{d \times d}$  a positive semi-definite preconditioning matrix.

**Random Walk Metropolis** The (pre-conditioned) random walk Metropolis (RWMH) [Metropolis et al. \(1953\)](#) with and step size  $h$  has proposal kernel

$$Q^{RW}(x, dy) = \mathcal{N}(dy \mid x, hG).$$

**Metropolis-adjusted Langevin Algorithm** In the (pre-conditioned) Metropolis-adjusted Langevin algorithm (MALA) ([Roberts and Tweedie, 1996](#); [Stramer and Roberts, 2007](#)) has proposal kernel

$$Q^L(x, dy) = \mathcal{N}\left(dy \mid x + \frac{h}{2}G\nabla \log \pi(x), hG\right).$$

**Barker Proposal** Let  $C$  denote the Cholesky factor of  $G$ , set  $\tau_i^2 = hG_{ii}$ , let  $\mu(\cdot)_{\tau_i}$  be the probability density function of  $\mathcal{N}(0, \tau_i^2)$ , and let  $c_i(x) = (\nabla \log \pi(x) \cdot C^T)_i$ . For a proposal state  $y \in \mathbb{R}^d$ , let  $z_i = ((C^\top)^{-1}(y - x))_i$ . The (pre-conditioned) Barker proposal [Livingstone and Zanella \(2022\)](#) for the  $i$ th coordinate is

$$Q_i^B(x, dy_i) = 2 \frac{\mu_{\tau_i}(y_i - x_i)}{1 + e^{-z_i c_i(x)}} dy_i,$$

and the full Barker kernel  $Q^B$  is

$$Q^B(x, dy) = \prod_{i=1}^d Q_i^B(x, dy_i).$$

**Hamiltonian Monte Carlo** Hamiltonian Monte Carlo (HMC) ([Neal, 2011](#)) is defined on an extended state space with the random momentum vector  $\eta \stackrel{\text{indep}}{\sim} \mathcal{N}(0, G)$ , yielding the joint density

$$\bar{\pi}(x, \eta) \propto \exp\left\{\log(\pi(x)) - \frac{1}{2}\eta^\top G \eta\right\}.$$

HMC updates a new state using the leapfrog integrator, which proceeds according to the updates

$$\begin{aligned} \eta(t + h/2) &= \eta(t) + (h/2)\nabla_x \log \pi(x(t)) \\ x(t + h) &= x(t) + hG\eta(t + h/2) \\ \eta(t + h) &= \eta(t + h/2) + (h/2)\nabla_x \log \pi(x(t + h)). \end{aligned}$$

Taking  $x(0) = x$ , the new proposal state is  $y = x(Lh)$  with  $\eta(0) \sim \mathcal{N}(0, G)$ .

## B PROOFS

For readability, we write the proposal kernel as  $K(x, h, dy) = K_{\log h}(x, dy)$ , so

$$X_j^{(t+1)} \sim K\left(X_j^{(t)}, h_N^{(t)}, \cdot\right).$$

Or, equivalently,

$$X_j^{(t+1)} = \begin{cases} Y_j^{(t)}, & \text{with probability } \alpha_j^{(t)} \\ X_j^{(t)}, & \text{with probability } 1 - \alpha_j^{(t)}, \end{cases}$$

with  $Y_j^{(t)}$  being the proposal state based on the proposal distribution  $Q_{h_N^{(t)}}(X_j^{(t)}, \cdot)$  and

$$\begin{aligned} \alpha_j^{(t)} &:= g(X_j^{(t)}, Y_j^{(t)}, h_N^{(t)}) \\ &:= \min \left\{ 1, \frac{\pi(Y_j^{(t)}) q_{h_N^{(t)}}(X_j^{(t)}, Y_j^{(t)})}{\pi(X_j^{(t)}) q_{h_N^{(t)}}(Y_j^{(t)}, X_j^{(t)})} \right\}, \end{aligned}$$

where  $q_h(x, \cdot)$  is the probability density of  $Q_h(x, \cdot)$  and  $h_N^{(t)}$  is the step size at time  $t$ . If we treat the proposed states  $Y_{1:N}^{(t)}$  as additional random variables, we can write

$$\left(X_j^{(t+1)}, Y_j^{(t+1)}\right) \sim T(X_j^{(t)}, Y_j^{(t)}, h_N^{(t)}, \cdot, \cdot)$$

for an appropriate Markov kernel  $T(x, y, h, dx', dy')$ . Now define the empirical distribution of  $N$  particles  $X_j^{(t)}$  at  $t$ th iteration as

$$\nu_N^{(t)} := \frac{1}{N} \sum_{j=1}^N \delta_{X_j^{(t)}}.$$

Similarly, we can define the empirical distribution of  $N$  particles  $(X_j^{(t)}, Y_j^{(t)})$  at the  $t$ th iteration as

$$\xi_N^{(t)} := \frac{1}{N} \sum_{i=1}^N \delta_{(X_i^{(t)}, Y_i^{(t)})}.$$

Taking the initial step size  $h_N^{(0)}$  as fixed, we can rewrite the  $t$ th step size  $h_N^{(t)}$  update equation given in Eq. (5) as

$$h_N^{(t)} = h_N^{(t-1)} \exp \left\{ \frac{1}{t} \left( \frac{1}{N} \sum_{j=1}^N \alpha_j^{(t-1)} - \bar{\alpha}_* \right) \right\}.$$

Letting  $\bar{h}^{(0)} = h_N^{(0)}$ ,  $\bar{\nu}^{(0)} = \hat{\pi}^{(0)}$  and  $\bar{\xi}^{(0)}(dx, dy) = Q_{h_N^{(0)}}(x, dy) \hat{\pi}^{(0)}(dx)$ , then we can recursively define  $\bar{h}^{(t)}$ ,  $\bar{\nu}^{(t)}$  and  $\bar{\xi}^{(t)}$  by

$$\begin{aligned} \bar{h}^{(t)} &:= \bar{h}^{(t-1)} \exp \left( \frac{1}{\sqrt{t}} \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle - \bar{\alpha}_* \right), \\ \bar{\nu}^{(t)} &:= \int K(x, \bar{h}^{(t-1)}, \cdot) \bar{\nu}^{(t-1)}(dx), \\ \bar{\xi}^{(t)} &:= \int T(x, y, \bar{h}^{(t-1)}, \cdot, \cdot) \bar{\xi}^{(t-1)}(dx, dy), \end{aligned}$$

where

$$\begin{aligned} \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle &:= \mathbb{E}_{(X, Y) \sim \bar{\xi}^{(t-1)}} \left[ g(X, Y, \bar{h}^{(t-1)}) \right] \\ &:= \int g(x, y, \bar{h}^{(t-1)}) \bar{\xi}^{(t-1)}(dx, dy). \end{aligned}$$

Let  $\mathcal{P}(S)$  denote the set of probability measures on the measurable space  $(S, \mathcal{B})$ . Let  $\Xi_N^{(t)}$  and  $\Pi_N^{(t)}$  denote the probability distributions of the respective empirical measures  $\xi_N^{(t)}$  and  $\nu_N^{(t)}$ , so  $\xi_N^{(t)} \in \mathcal{P}(\mathcal{P}(\mathbb{R}^d))$  and  $\Pi_N^{(t)} \in \mathcal{P}(\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d))$  respectively.

We will make use of the following two results.

**Theorem B.1.** (*Sznitman, 1991*)  $\{P_N^{(t)}\}$  is  $\bar{\nu}^{(t)}$ -chaotic if and only if

$$P_N^{(t)} \circ (\nu_N^{(t)})^{-1} \xrightarrow{d} \delta(\bar{\nu}^{(t)})$$

in  $\mathcal{P}(\mathcal{P}(S))$ .

**Theorem B.2.** If  $X, X_n : \Omega \rightarrow \mathbb{R}^d$  are random variables such that  $X_n \xrightarrow{P} X$ , then  $f(X_n) \xrightarrow{L^1} f(X)$  for all  $f \in C_b(\mathbb{R}^d)$ .

### B.1 Proof of Theorem 3.3

Theorem 3.3 will follow from a series of propositions and lemmas, which we prove in the subsequent subsections.

**Proposition B.3.** Under Assumption 3.4, if  $(X_{1:N}^{(t-1)}, Y_{1:N}^{(t-1)})$  is  $\bar{\xi}^{(t-1)}$ -chaotic and  $h_N^{(t-1)} \xrightarrow{P} \bar{h}^{(t-1)}$ , then  $(X_{1:N}^{(t)}, Y_{1:N}^{(t)})$  is  $\bar{\xi}^{(t)}$ -chaotic.

**Proposition B.4.** Under Assumptions 3.4.(a) and 3.4.(b), if  $(X_{1:N}^{(t-1)}, Y_{1:N}^{(t-1)})$  is  $\bar{\xi}^{(t-1)}$ -chaotic and  $h_N^{(t-1)} \xrightarrow{P} \bar{h}^{(t-1)}$ , then  $h_N^{(t)} \xrightarrow{P} \bar{h}^{(t)}$ .

Note that  $\{(X_j^{(0)}, Y_j^{(0)})\}_{j=1}^N$  is  $\bar{\xi}^{(0)}$ -chaotic since the samples are i.i.d. and  $h_N^{(0)} = \bar{h}^{(0)}$  is fixed. By induction, it follows from Proposition B.3 and Proposition B.4 that for any  $t \in \mathbb{N}$ ,  $\{(X_j^{(t)}, Y_j^{(t)})\}_{j=1}^N$  is  $\bar{\xi}^{(t)}$ -chaotic and  $h_N^{(t)} \xrightarrow{P} \bar{h}^{(t)}$ .

The following lemma implies that there is subsequence  $\{\Pi_{N_l}^{(t)}\}_l$  that weakly converges.

**Lemma B.5.** Under Assumption 3.4.(c), for any fixed  $t \in \mathbb{N}$ ,  $\{\Pi_N^{(t)}\}_N$  is tight and relatively compact.

**Lemma B.6.** Under Assumption 3.4, if  $X_{1:N}^{(t-1)}$  is  $\bar{\nu}_{t-1}$ -chaotic,  $h_N^{(t-1)} \xrightarrow{P} \bar{h}^{(t-1)}$ , and  $\{\Pi_{N_l}^{(t)}\}_l$  is a convergent subsequence of  $\{\Pi_N^{(t)}\}_N$  with limit  $\Pi^{(t)}$ , then  $\Pi^{(t)}$  is a Dirac measure concentrated on  $\bar{\nu}^{(t)}$ .

**Lemma B.7.** Under assumptions and conditions of Lemma B.6,  $X_{1:N}^{(t)}$  is  $\bar{\nu}^{(t)}$ -chaotic.

We also know that  $X_{1:N}^{(0)}$  is  $\bar{\nu}^{(0)}$ -chaotic (due to the fact that the  $X_{1:N}^{(0)}$  are i.i.d.), and  $h_N^{(0)} \xrightarrow{P} \bar{h}^{(0)}$  for any  $t \in \mathbb{N}$ . Hence, by recursively using Lemma B.7, we conclude that  $X_{1:N}^{(t)}$  is  $\bar{\nu}^{(t)}$ -chaotic for any  $t \in \mathbb{N}$ .

### B.2 Proof of Proposition B.3

To prove Proposition B.3, we will first show that the probability distribution  $\{\Xi_N^{(t)}\}_N$  of empirical measure  $\{\xi_N^{(t)}\}_N$  is tight and relatively compact, which guarantees that there exists convergent subsequence  $\{\Xi_{N_l}^{(t)}\}_l$ . Then we will prove that every convergent subsequence  $\{\Xi_{N_l}^{(t)}\}_l$  has the same limit  $\Xi^{(t)}$  that is a Dirac measure concentrated on  $\bar{\xi}^{(t)}$ .

**Lemma B.8.** Under Assumption 3.4.(c), for any fixed  $t \in \mathbb{N}$ ,  $\{\Xi_N^{(t)}\}_N$  is tight and relatively compact.

**Lemma B.9.** Under Assumption 3.4, suppose  $(X_{1:N}^{(t-1)}, Y_{1:N}^{(t-1)})$  is  $\bar{\xi}^{(t-1)}$ -chaotic and  $h_N^{(t-1)} \xrightarrow{P} \bar{h}^{(t-1)}$ . Let  $\{\Xi_{N_l}^{(t)}\}_l$  be a convergent subsequence of  $\{\Xi_N^{(t)}\}_N$  with limit  $\Xi^{(t)}$ . Then  $\Xi^{(t)}$  is a Dirac measure concentrated on  $\bar{\xi}^{(t)}$ .

The proofs of Lemma B.8 and B.9 are deferred to Appendix B.2.1 and B.2.2. Relative compactness of  $\{\Xi_N^{(t)}\}_N$  implies that there is subsequence  $\{\Xi_{N_l}^{(t)}\}_l$  which weakly converges. In this proof, with a slight abuse of notation, we use  $\{\Xi_N^{(t)}\}_N$  to represent the convergent subsequence  $\{\Xi_{N_l}^{(t)}\}_l$ . By Lemma B.8 and Lemma B.9, we have

$$\xi_N^{(t)} \xrightarrow{d} \bar{\xi}^{(t)}.$$

Theorem B.1 guarantees that  $(X_{1:N}^{(t)}, Y_{1:N}^{(t)})$  is  $\bar{\xi}^{(t)}$ -chaotic.

### B.2.1 Proof of Lemma B.8

The proof is similar to that of Lemma 2.2 in Sirignano and Spiliopoulos (2020). For a fixed  $t \in \mathbb{N}$ ,  $\forall \epsilon > 0$ , we want to show that there is a compact subset  $K$  of  $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$  such that

$$\sup_{N \in \mathbb{N}} \mathbb{P} \left[ \xi_N^{(t)} \notin K \right] < \epsilon.$$

For each  $L > 0$ , define  $K_L = [-L, L]^{2d}$ , and by Assumption 3.4.(c), there exists  $C$  such that

$$\begin{aligned} \mathbb{E} \left[ \xi_N^{(t)} (\mathbb{R}^{2d} \setminus K_L) \right] &\leq \frac{1}{N} \sum_{j=1}^N \mathbb{P} \left[ \left\| (X_j^{(t)}, Y_j^{(t)}) \right\| \geq L \right] \\ &= \mathbb{P} \left[ \left\| (X_j^{(t)}, Y_j^{(t)}) \right\| \geq L \right] \\ &\leq \frac{\mathbb{E} \left[ \left\| (X_j^{(t)}, Y_j^{(t)}) \right\| \right]}{L} \\ &\leq \frac{\mathbb{E} \left[ \left\| X_j^{(t)} \right\| + \left\| Y_j^{(t)} \right\| \right]}{L} \\ &= \frac{C}{L}. \end{aligned}$$

For  $L > 0$ , define the compact set

$$\hat{K}_L = \left\{ \nu : \nu (\mathbb{R}^{2d} \setminus K_{(L+j)^2}) < \frac{1}{\sqrt{L+j}} \text{ for all } j \in \mathbb{N} \right\}.$$

Observe that

$$\begin{aligned} \mathbb{P} \left[ \xi_N^{(t)} \notin \hat{K}_L \right] &\leq \sum_{j=1}^{\infty} \mathbb{P} \left[ \xi_N^{(t)} (\mathbb{R}^{2d} \setminus K_{(L+j)^2}) > \frac{1}{\sqrt{L+j}} \right] \leq \sum_{j=1}^{\infty} \frac{\mathbb{E} \left[ \xi_N^{(t)} (\mathbb{R}^{2d} \setminus K_{(L+j)^2}) \right]}{1/\sqrt{L+j}} \\ &\leq \sum_{j=1}^{\infty} \frac{C}{(L+j)^2/\sqrt{L+j}} \leq \sum_{j=1}^{\infty} \frac{C}{(L+j)^{3/2}}, \end{aligned}$$

so  $\lim_{L \rightarrow \infty} \mathbb{P}[\xi_N^{(t)} \notin \hat{K}_L] = 0$ . Thus, the sequence  $(\xi_N^{(t)})_{N \in \mathbb{N}}$  is tight, which implies it is relatively compact.

### B.2.2 Proof of Lemma B.9

For each  $f \in C_b(\mathbb{R}^d \times \mathbb{R}^d)$ , define the map  $F : \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) \rightarrow \mathbb{R}_+$  by

$$F(\xi) = \left| \langle f, \xi \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right|.$$

We would like to show that  $\lim_{N \rightarrow \infty} \mathbb{E}_{\xi \sim \Xi_N^{(t)}} [F(\xi)] = 0$ . Together with the fact that subsequence  $\{\Xi_N^{(t)}\}_N$  converges in distribution to  $\Xi^{(t)}$ , we would have  $\mathbb{E}_{\xi \sim \Xi^{(t)}} [F(\xi)] = 0$ . Then we can conclude that  $\Xi^{(t)}$  is a Dirac measure concentrated on  $\bar{\xi}^{(t)}$ .

Letting  $\langle f, \tilde{\xi}_N^{(t)} \rangle = \frac{1}{N} \sum_{j=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot) \rangle$ , we have

$$\begin{aligned} \mathbb{E}_{\xi \sim \Xi_N^{(t)}} [F(\xi)] &= \mathbb{E} \left[ F \left( \xi_N^{(t)} \right) \right] \\ &= \mathbb{E} \left[ \left| \langle f, \xi_N^{(t)} \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right| \right] \\ &= \mathbb{E} \left[ \left| \langle f, \xi_N^{(t)} \rangle - \langle f, \tilde{\xi}_N^{(t)} \rangle + \langle f, \tilde{\xi}_N^{(t)} \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right| \right] \\ &\leq \mathbb{E} \left[ \left| \langle f, \xi_N^{(t)} \rangle - \langle f, \tilde{\xi}_N^{(t)} \rangle \right| \right] + \mathbb{E} \left[ \left| \langle f, \tilde{\xi}_N^{(t)} \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right| \right], \end{aligned}$$

We will prove that  $\lim_{N \rightarrow \infty} \mathbb{E}_{\xi \sim \Xi_N^{(t)}} [F(\xi)] = 0$  by the following two lemmas:

**Lemma B.10.** *Under the assumptions of Lemma B.9, we have*

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \langle f, \xi_N^{(t)} \rangle - \langle f, \tilde{\xi}_N^{(t)} \rangle \right| \right] = 0.$$

**Lemma B.11.** *Under the assumptions of Lemma B.9, we have*

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \langle f, \tilde{\xi}_N^{(t)} \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right| \right] = 0.$$

By Lemma B.10 and B.11, we have

$$\lim_{N \rightarrow \infty} \mathbb{E}_{\Xi_N^{(t)}} [F(\xi)] = 0.$$

Since  $F$  is uniformly bounded, we have

$$\lim_{N \rightarrow \infty} \mathbb{E}_{\Xi_N^{(t)}} [F(\xi)] = \mathbb{E}_{\Xi^{(t)}} [F(\xi)] = 0.$$

Since this holds true for any  $f \in C_b(\mathbb{R}^d \times \mathbb{R}^d)$ , using the fact that  $F$  is nonnegative and by Theorem 5.1 from Billingsley (2013), we have that  $\Xi_N^{(t)}$  has a limit point that is a Dirac measure concentrated on  $\bar{\xi}^{(t)}$ .

### B.2.3 Proof of Lemma B.10

Let  $\mathcal{F}_N^{(t)}$  be the  $\sigma$ -algebra generated by  $X_{1:N}^{(1:m)}$  and  $Y_{1:N}^{(1:m)}$ , for  $i = 1, 2, \dots, N$  and  $m = 1, 2, \dots, t$ . Note that

$$\begin{aligned} & \mathbb{E} \left[ \left| \langle f, \xi_N^{(t)} \rangle - \langle f, \tilde{\xi}_N^{(t)} \rangle \right| \mid \mathcal{F}_N^{(t-1)} \right] \\ &= \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f(X_j^{(t)}, Y_j^{(t)}) - \frac{1}{N} \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot) \rangle \right| \mid \mathcal{F}_N^{(t-1)} \right] \\ &\leq \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f(X_j^{(t)}, Y_j^{(t)}) - \frac{1}{N} \sum_{i=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot) \rangle \right| \mid \mathcal{F}_N^{(t-1)} \right] \\ &+ \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot) \rangle - \frac{1}{N} \sum_{i=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot) \rangle \right| \mid \mathcal{F}_N^{(t-1)} \right] \\ &= \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f(X_j^{(t)}, Y_j^{(t)}) - \frac{1}{N} \sum_{i=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot) \rangle \right| \mid \mathcal{F}_N^{(t-1)} \right] \\ &+ \left| \frac{1}{N} \sum_{j=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot) \rangle - \frac{1}{N} \sum_{j=1}^N \langle f, T(X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot) \rangle \right|. \end{aligned}$$

Since  $f \in C_b(\mathbb{R}^d \times \mathbb{R}^d)$ , there exists  $M$  such that for any  $i = 1, 2, \dots, N$ , we have

$$\text{Var} \left[ f(X_j^{(t)}, Y_j^{(t)}) \mid \mathcal{F}_N^{(t-1)} \right] \leq M < \infty.$$



It follows that

$$\begin{aligned}
 & \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle \right| \middle| \mathcal{F}_N^{(t-1)} \right] \\
 &= \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \mathbb{E} \left[ \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) \middle| \mathcal{F}_N^{(t-1)} \right] \right| \middle| \mathcal{F}_N^{(t-1)} \right] \\
 &\leq \mathbb{E}^{1/2} \left[ \left( \frac{1}{N} \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \mathbb{E} \left[ \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) \middle| \mathcal{F}_N^{(t-1)} \right] \right)^2 \middle| \mathcal{F}_N^{(t-1)} \right] \\
 &= \text{Var}^{1/2} \left[ \frac{1}{N} \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) \middle| \mathcal{F}_N^{(t-1)} \right] \\
 &\leq \left( \frac{M}{N} \right)^{1/2}.
 \end{aligned}$$

The last inequality holds true because  $(X_j^{(t)}, Y_j^{(t)})$  are independent given  $\mathcal{F}_N^{(t-1)}$ . Therefore, we have

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle \right| \middle| \mathcal{F}_N^{(t-1)} \right] = 0.$$

And  $\mathbb{E} \left[ \left| \frac{1}{N} \sum_{i=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \sum_{i=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle \right| \middle| \mathcal{F}_N^{(t-1)} \right]$  is bounded due to the fact that  $f \in C_b(\mathbb{R}^d \times \mathbb{R}^d)$ , then by dominated convergence theorem, we have

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \frac{1}{N} \sum_{i=1}^N f \left( X_j^{(t)}, Y_j^{(t)} \right) - \frac{1}{N} \sum_{i=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle \right| \right] = 0.$$

Note that

$$\begin{aligned}
 & \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle - \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle \right| \right] \\
 &\leq \frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[ \left| \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle \right| \right] \\
 &= \mathbb{E} \left[ \left| \langle f, T \left( X_1^{(t-1)}, Y_1^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, T \left( X_1^{(t-1)}, Y_1^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle \right| \right].
 \end{aligned}$$

Since we know  $(X_1^{(t-1)}, Y_1^{(t-1)}, h_N^{(t-1)}) \xrightarrow{p} (X_1^{(t-1)}, Y_1^{(t-1)}, \bar{h}^{(t-1)})$  and  $\langle f, T(x, y, h, \cdot, \cdot) \rangle$  is bounded and continuous with respect to  $(x, y, h)$ , thanks to Theorem B.2, we have

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}, \cdot, \cdot \right) \rangle - \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle \right| \right] = 0.$$

### B.2.4 Proof of Lemma B.11

Note that

$$\begin{aligned} & \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right|^2 \right] \\ &= \frac{1}{N} \mathbb{E} \left[ \left| \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right|^2 \right] \\ &+ \frac{N^2 - N}{N^2} \mathbb{E} \left[ \left( \langle f, T \left( X_i^{(t-1)}, Y_i^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right) \left( \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right) \right]. \end{aligned}$$

The above expectation is taken with respect to  $(X_j^{(t-1)}, Y_{i,t-1})$ . For any fixed  $(X_j^{(t-1)}, Y_{i,t-1})$

$$\begin{aligned} \langle f, T \left( X_j^{(t-1)}, Y_{1,t-1}, \bar{h}^{(t-1)}, \cdot \right) \rangle &= \int f(s_1, s_2) T \left( X_j^{(t-1)}, Y_{1,t-1}, \bar{h}^{(t-1)}, ds_1, ds_2 \right) \\ &= \int f \left( X_j^{(t-1)}, s_2 \right) \left( 1 - g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) \right) Q_{\bar{h}^{(t-1)}} \left( X_j^{(t-1)}, ds_2 \right) \\ &+ \int f \left( Y_j^{(t-1)}, s_2 \right) g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) Q_{\bar{h}^{(t-1)}} \left( Y_j^{(t-1)}, ds_2 \right). \end{aligned}$$

By Assumption 3.4.(a) and 3.4.(b), it is bounded and continuous with respect to  $(X_j^{(t-1)}, Y_j^{(t-1)})$ . Then thanks to that fact that  $(X_{1:N}^{(t-1)}, Y_{1:N}^{(t-1)})$  is  $\bar{\xi}^{(t-1)}$ -chaotic, we have

$$\begin{aligned} & \lim_{N \rightarrow \infty} \mathbb{E} \left[ \left( \langle f, T \left( X_i^{(t-1)}, Y_i^{(t-1)}, \bar{h}^{(t-1)}, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right) \left( \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right) \right] \\ &= \left( \int \langle f, T \left( s_1, s_2, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle d\bar{\xi}^{(t-1)}(ds_1, ds_2) \right)^2 = 0. \end{aligned}$$

Since  $f \in C_b(\mathbb{R}^d \times \mathbb{R}^d)$ , we have  $\mathbb{E} \left[ \left| \langle f, T \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)}, \cdot, \cdot \right) \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right|^2 \right] < \infty$ . Therefore, we have

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \langle f, \tilde{\xi}_N^{(t)} \rangle - \langle f, \bar{\xi}^{(t)} \rangle \right| \right] = 0.$$

### B.3 Proof of Proposition B.4

Recall that the step size update is

$$\begin{aligned} h_N^{(t)} &= h_N^{(t-1)} \exp \left( \frac{1}{t} \left( \left( \frac{1}{N} \sum_{j=1}^N \alpha_j^{(t-1)} \right) - \bar{\alpha}_* \right) \right) \\ &= h_N^{(t-1)} \exp \left( \frac{1}{t} \left( \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - \bar{\alpha}_* \right) \right). \end{aligned}$$

Note that

$$\begin{aligned} & \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle \\ &= \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) \\ &+ \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) - \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle. \end{aligned} \tag{9}$$

The first term in Eq. (9) can be rewritten as

$$\begin{aligned} & \mathbb{E} \left[ \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) \right] \\ &= \mathbb{E} \left[ g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) \right]. \end{aligned}$$

Since  $(X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)}) \xrightarrow{p} (X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)})$  and  $g$  is bounded and continuous, it follows from Theorem B.2 that

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[ \left| \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) - \frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) \right| \right] = 0. \quad (10)$$

The second term of Eq. (9) can be rewritten as

$$\frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, \bar{h}^{(t-1)} \right) - \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle = \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \xi_N^{(t-1)} \rangle - \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle.$$

Since  $(X_{1:N}^{(t-1)}, Y_{1:N}^{(t-1)})$  is  $\bar{\xi}^{(t-1)}$ -chaotic, by Theorem B.1 we have

$$\xi_N^{(t-1)} \xrightarrow{d} \bar{\xi}^{(t-1)}.$$

By Assumptions 3.4.(a) and 3.4.(b), we have  $g(x, y, \bar{h}^{(t-1)})$  is continuous with respect to  $(x, y)$ , thus

$$\langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \xi_N^{(t-1)} \rangle \xrightarrow{p} \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle. \quad (11)$$

It follows from Eq. (10) and Eq. (11) that

$$\frac{1}{N} \sum_{j=1}^N g \left( X_j^{(t-1)}, Y_j^{(t-1)}, h_N^{(t-1)} \right) \xrightarrow{p} \langle g(\cdot, \cdot, \bar{h}^{(t-1)}), \bar{\xi}^{(t-1)} \rangle$$

and therefore  $h_N^{(t)} \xrightarrow{p} \bar{h}^{(t)}$ .

#### B.4 Proofs of Lemmas B.5 to B.7.

The proof of Lemma B.5 is similar to that of Lemma B.8; the proof of Lemma B.6 is similar to that of Lemma B.9; the proof of Lemma B.7 is similar to that of Proposition B.3.

## C ADDITIONAL RESULTS AND EXPERIMENTS

### C.1 Ablation Study On $T$

In Section 3.3, we discussed how to choose length of Markov chains properly: Eq. (7) and Eq. (8) guarantee that we can obtain efficient samples under different dimension  $d$  without too large computation cost. Now, we would like to explore how the lower bounds would change under different number of iterations  $T$ .

Figures 10 to 12 display the evolution of several lower bounds for  $X_1$  of Gaussian model (Section 4.1),  $X_1$  of Neal-funnel shape model (Section 4.2) and  $\lambda_{61}$  of cancer classification using horseshoe prior (Section 4.4) respectively, which shows that when  $\rho$  is large (greater than 0.1), the proposed lower bounds are quite underestimated. These results also confirm the lower bounds become nearly stationary at our proposed number of iterations  $T$ .

### C.2 Additional Results for Neal-Funnel Shape Model

In Section 4.2, we validate TADDAA on Neal's funnel distribution approximated by mean-field Gaussian VI. Fig. 13 provides reliability checks for the Neal's funnel experiment Fig. 4.




Figure 10: Ablation study for  $T$  on  $X_1$  of Gaussian model with correlated coordinates,  $d = 30$ .




Figure 11: Ablation study for  $T$  on  $X_1$  of Neal-funnel shape model,  $d = 30$ .




Figure 12: Ablation study for  $T$  on  $\lambda_{61}$  of Neal-funnel shape model,  $d = 100$ .




Figure 13: Reliability checks on TADDAA using Barker applied to the mean-field VI approximation to Neal’s funnel distribution of varying dimensionality.

Next, we will validate the TADDAA on Neal’s funnel distribution approximated by mean-field t-distribution VI. The lower bounds  $B_{\text{mean}}$  and  $B_{\text{tail}}$  are computed according to Eq. (1) using functionals of interest  $\mathcal{F}_{0.5}$  and  $\mathcal{F}_{0.9}$  defined based on Eq. (4).

Fig. 14 shows that the median and tail quantile lower bounds constructed using Barker are quite precise and remain valid when  $d$  and  $hat{k}$  are large. As we can see,  $\hat{k}$  is large ( $\hat{k} > 0.7$ ) for almost any dimension, indicating VI approximation is not reliable. But if we are only interested in whether median of  $X_2$ , mean field t-distribution VI should be considered as reliable.

### C.3 Reliability Check for Bayesian Neural Network

In Section 4.5, we validate TADDAA on Bayesian neural network approximated by mean-field Gaussian VI. The reliability check is shown in Fig. 15: due to the fact that the variance is dramatically underestimated, only Barker chains pass the reliability check, which is also consistent with the fact that Barker kernel provides the best lower bounds as displayed in Fig. 9.

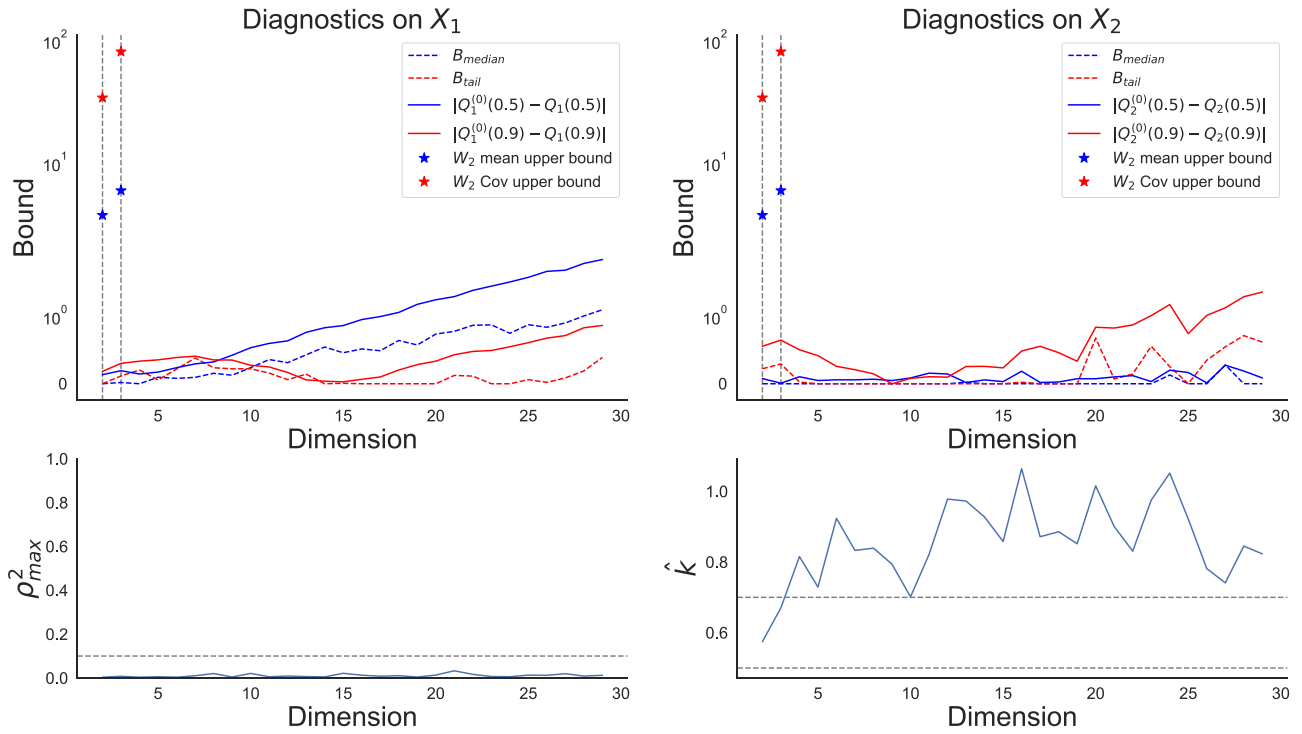


Figure 14: Diagnostics for Neal-funnel shape model, where TADDAA uses the Barker proposal.  $Q_1(0.5)$  and  $Q_1(0.9)$  denote the median and 90%-quantile of  $X_1$ ,  $Q_2(0.5)$  and  $Q_2(0.9)$  denote the median and 90%-quantile of  $X_2$ .




Figure 15: Reliability checks for TADDAA using different kernels applied to Bayesian neural network approximated by mean-field Gaussian VI for candy power ranking dataset.