

---

# Coreset Markov chain Monte Carlo

---

**Naitong Chen**

Department of Statistics  
University of British Columbia  
naitong.chen@stat.ubc.ca

**Trevor Campbell**

Department of Statistics  
University of British Columbia  
trevor@stat.ubc.ca

## Abstract

A Bayesian coreset is a small, weighted subset of data that replaces the full dataset during inference in order to reduce computational cost. However, state of the art methods for tuning coreset weights are expensive, require nontrivial user input, and impose constraints on the model. In this work, we propose a new method—*Coreset MCMC*—that simulates a Markov chain targeting the coreset posterior, while simultaneously updating the coreset weights using those same draws. Coreset MCMC is simple to implement and tune, and can be used with any existing MCMC kernel. We analyze Coreset MCMC in a representative setting to obtain key insights about the convergence behaviour of the method. Empirical results demonstrate that Coreset MCMC provides higher quality posterior approximations and reduced computational cost compared with other coreset construction methods. Further, compared with other general subsampling MCMC methods, we find that Coreset MCMC has a higher sampling efficiency with competitively accurate posterior approximations.

## 1 INTRODUCTION

Bayesian inference provides a flexible and principled framework for parameter estimation and uncertainty quantification when working with complex statistical models. Markov chain Monte Carlo (MCMC) (Robert and Casella, 2004; Robert and Casella, 2011; Gelman et al., 2013, Ch. 11,12) is the standard methodology for conducting Bayesian inference, and involves simulating

a Markov chain whose stationary distribution is the Bayesian posterior. However, in the large-scale data setting, standard MCMC methods become costly, as simulating each Markov transition involves iterating over the entire data set; and typically, many steps are needed to obtain reasonable estimates of posterior expectations.

To reduce the computational cost incurred by large datasets, *subsampling MCMC* methods repeatedly subsample the data and simulate the next state of the chain based on only that subsample, instead of the full dataset (Banterle et al., 2019; Quiroz et al., 2019; Maclaurin and Adams, 2014; Korattikara et al., 2014; Bardenet et al., 2014; Welling and Teh, 2011; Chen et al., 2014). Some of these methods target the exact posterior distribution asymptotically, but require a full pass over the data per accepted step (Banterle et al., 2019), require model-specific design of log-likelihood surrogates (Maclaurin and Adams, 2014), or have progressively slower convergence (Welling and Teh, 2011). Other methods approximate various aspects of the transition with a subsample—e.g., the accept-reject decision (Korattikara et al., 2014; Bardenet et al., 2014), or dynamics that generate a proposal (Chen et al., 2014; Baker et al., 2019)—but the performance of such methods is dependent on the existence of good control variates to keep (gradient) log-likelihood estimate variance low (Quiroz et al., 2019). Control variate design is model-specific in general, while automated methods impose limitations on the model (e.g., Taylor expansion-based control variates require continuous latent variables). For more in-depth reviews of subsampling MCMC methods, see Bardenet et al. (2017); Quiroz et al. (2018); Nemeth and Fearnhead (2021).

Another approach to handling large-scale Bayesian inference problems are *Bayesian coresets* (Huggins et al., 2016; Campbell and Broderick, 2019, 2018; Campbell and Beronov, 2019; Manousakas et al., 2020; Naik et al., 2022; Chen et al., 2022), which consist of a fixed, weighted subsample of the data that replaces the full dataset during inference. The intuition is that there is typically redundancy present in large data, and so it

---

Proceedings of the 27<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

should be possible to capture the information needed for inference in a single small subsample. Indeed, recent work has shown that a coreset of size  $O(\log N)$  suffices to provide a near-exact approximation of a posterior for  $N$  data points in a wide class of models (Naik et al., 2022, Thm. 4.1,4.2; Chen et al., 2022). Coresets do not require bespoke kernel or control-variate design; once a coreset is constructed, any generic MCMC kernel can be applied. Coresets also typically preserve important structure (e.g., unidentifiability, heavy tails) as they are built using the original model’s log-likelihood terms.

Constructing a good coreset efficiently remains a challenge. Early methods based on importance sampling (Huggins et al., 2016) are simple and computationally efficient, but generally produce unreliable approximations in practice. Methods based on sparse regression with a finite-dimensional log-likelihood projection (Campbell and Broderick, 2019, 2018; Zhang et al., 2021) are sensitive to the projection and require significant expert user tuning effort. Sequential greedy KL minimization approaches (Campbell and Beronov, 2019) are able to produce high quality coresets, but involve a slow and difficult-to-tune inner-loop weight optimization with outer-loop point selection. Two recent methods—Sparse Hamiltonian Flows (SHF) and Quasi-Newton Coresets (QNC) (Naik et al., 2022; Chen et al., 2022)—instead quickly select the points in the coreset via uniform subsampling, and then run a single joint weight optimization. SHF avoids inner-outer loop procedures entirely, but is limited to differentiable log-posterior densities and does not come with a convergence guarantee on the weights. QNC applies more broadly and has a guaranteed weight convergence, but only if certain expectations can be evaluated exactly; in practice, one requires inner-loop MCMC. Both methods tend to require significant user expertise and effort when tuning.

This work introduces *Coreset MCMC*, a new construction method that is simpler and faster to implement and tune compared with previous coreset methods. Coreset MCMC can be thought of as a meta-algorithm that wraps an existing MCMC kernel: the method iterates between (1) taking a step with the kernel targeting the coreset posterior, and (2) adapting the coreset weights using the current draw, and is related to both adaptive stochastic approximation algorithms (see Benveniste et al., 1990, pp. 31-33) and adaptive MCMC (see Andrieu and Thoms, 2008). We show that when the optimal coreset is exact, Coreset MCMC will produce a coreset that converges to the exact posterior, and hence is an asymptotically exact method. We also analyze Coreset MCMC using a representative model to obtain key insights in tuning the method. This paper concludes with experiments demonstrating

that coreset MCMC provides higher quality posterior approximations than other coreset methods, and improved sampling efficiency compared with subsampling MCMC methods.

## 2 BACKGROUND

We are given a dataset  $(X_n)_{n=1}^N$  of  $N$  observations, a log-likelihood  $\ell_n(\theta) := \log p(X_n | \theta)$  for observation  $n$  given  $\theta \in \Theta$ , and a prior density  $\pi_0(\theta)$ . The goal is to sample from the Bayesian posterior with density

$$\pi(\theta) := \frac{1}{Z} \exp\left(\sum_{n=1}^N \ell_n(\theta)\right) \pi_0(\theta),$$

where  $Z$  is the normalizing constant. The Bayesian coresets approach to reducing this cost involves replacing the sum of log-likelihoods for the full dataset with a small, weighted subsample (without loss of generality, we assume these are the first  $M$  points  $1, \dots, M$ ):

$$\pi_w(\theta) := \frac{1}{Z(w)} \exp\left(\sum_{m=1}^M w_m \ell_m(\theta)\right) \pi_0(\theta),$$

where  $w \in \mathbb{R}^M$ ,  $w \geq 0$  is a vector of nonnegative weights. In this work, we follow the setting in Naik et al. (2022) and Chen et al. (2022), where the coreset points are uniformly subsampled. If we can construct a set of weights  $w$  such that  $M \ll N$  and  $\pi_w \approx \pi$ , then we can generate draws using MCMC targeting  $\pi_w$  as a surrogate for draws from  $\pi$  at a low per-iteration cost.

Most recent coreset construction methods formulate the task as a variational inference problem, following Campbell and Beronov (2019):

$$w^* = \arg \min_{w \in \mathbb{R}^M} D_{\text{KL}}(\pi_w || \pi) \quad \text{s.t.} \quad w \in \mathcal{W}. \quad (1)$$

In this work, we assume that the feasible region  $\mathcal{W} \subseteq \mathbb{R}^M$  is convex and closed. We also assume that for all  $w \in \mathcal{W}$ ,  $Z(w) < \infty$ . In practice,  $\mathcal{W}$  is usually the set of nonnegative weights  $\mathbb{R}_+^M$ , but often also includes other constraints, e.g.,  $\sum_m w_m = N$ . This variational problem is also slightly unusual, in the sense that we cannot estimate the KL divergence *even up to a constant* because of the unknown normalization  $Z(w)$  that depends on the coreset weights  $w$ . However, we can write the  $M$ -dimensional KL gradient as an expectation under  $\pi_w$  that does not explicitly involve the normalization:

$$\begin{aligned} & \nabla_w D_{\text{KL}}(\pi_w || \pi) \\ &= \text{Cov}_{\pi_w} \left( \begin{bmatrix} \ell_1(\theta) \\ \vdots \\ \ell_M(\theta) \end{bmatrix}, \sum_m w_m \ell_m(\theta) - \sum_n \ell_n(\theta) \right). \quad (2) \end{aligned}$$

If we had access to i.i.d. draws from  $\pi_w$ , we could obtain an unbiased estimate of the gradient in Eq. (2) for stochastic gradient descent (Robbins and Monro, 1951; Bottou, 2004) by estimating the covariance using these i.i.d. draws, and estimating the full-data sum  $\sum_n \ell_n(\theta)$  with a subsample. Note that crucially, this estimate can be computed using only black-box evaluations of the log-likelihood, which makes this approach applicable to a wide range of models (e.g., with discrete variables). However, i.i.d. draws from  $\pi_w$  are usually not available, and one must resort to MCMC. This introduces an inner loop MCMC estimation step, which is slow and may need constant tuning as the coreset weights  $w$  evolve (Campbell and Beronov, 2019; Naik et al., 2022). In the next section, we draw inspiration from adaptive MCMC and adaptive stochastic approximation algorithms to develop a coreset construction method that more naturally interleaves MCMC and gradient updates.

### 3 CORESET MCMC

#### 3.1 Setup

Suppose at iteration  $t \in \mathbb{N}$ , for a fixed set of coreset weights  $w_t \in \mathcal{W}$ , we had access to  $K \geq 2$  i.i.d. draws  $\theta_t = (\theta_{t1}, \dots, \theta_{tK}) \in \Theta^K$  from  $\pi_{w_t}$ , and a subsample  $\mathcal{S}_t \subseteq [N]$  of  $S$  data points drawn uniformly without replacement from the full dataset. Then an unbiased estimate of the gradient in Eq. (2) is given by

$$g(w_t, \theta_t, \mathcal{S}_t) = \frac{1}{K-1} \sum_{k=1}^K \begin{bmatrix} \bar{\ell}_1(\theta_{tk}) \\ \vdots \\ \bar{\ell}_M(\theta_{tk}) \end{bmatrix} \left( \sum_m w_{tm} \bar{\ell}_m(\theta_{tk}) - \frac{N}{S} \sum_{s \in \mathcal{S}_t} \bar{\ell}_s(\theta_{tk}) \right), \quad (3)$$

where  $\bar{\ell}_n(\theta_{tk})$  are the centered log-likelihoods

$$\bar{\ell}_n(\theta_{tk}) = \ell_n(\theta_{tk}) - \frac{1}{K} \sum_{j=1}^K \ell_n(\theta_{tj}).$$

Since the estimate  $g(w_t, \theta_t, \mathcal{S}_t)$  is unbiased, i.e.,  $\mathbb{E}g(w_t, \theta_t, \mathcal{S}_t) = \nabla D_{\text{KL}}(\pi_{w_t} || \pi)$ , it can be used in a stochastic optimization scheme:<sup>1</sup>

$$w_{t+1} = \text{proj}_{\mathcal{W}}(w_t - \gamma_t g(w_t, \theta_t, \mathcal{S}_t)), \quad (4)$$

where  $\gamma_t > 0$  is a monotone decreasing learning rate sequence, and  $\text{proj}_{\mathcal{W}}$  projects the result onto the feasible set  $\mathcal{W}$ . However, in realistic scenarios, we do not have access to independent draws from  $\pi_w$ ; we are forced

<sup>1</sup>In Eqs. (4) and (5) we show standard projected stochastic gradient descent, but other methods such as ADAM (Kingma and Ba, 2014) and AdaGrad (Duchi et al., 2011) are also possible. We use ADAM in our experiments.

---

#### Algorithm 1 One iteration of Coreset MCMC

---

**Require:**  $w_t, \theta_t, \gamma_t, \kappa_w, S$   
 $\triangleright$  Subsample the data  
 $\mathcal{S}_t \leftarrow \text{Unif}(S, [N])$  (without replacement)  
 $\triangleright$  Compute gradient estimate  
 $g_t \leftarrow \text{Eq. (3)}$   
 $\triangleright$  Stochastic gradient step  
 $w_{t+1} \leftarrow w_t - \gamma_t g_t$   
 $\triangleright$  Step each Markov chain  
**for**  $k = 1, \dots, K$  **do**  
 $\theta_{(t+1)k} \sim \kappa_{w_{t+1}}(\cdot | \theta_{tk})$   
**end for**  
**return**  $w_{t+1}, \theta_{t+1}$

---

to take approximate draws using Markov chains. Our algorithm in Section 3.2 is based on using the stochastic optimization scheme in Eq. (4) but with  $K$  Markov chains instead of independent sequences.

#### 3.2 Algorithm

**Iteration:** Let  $\kappa_w$  be a family of Markov kernels parametrized by coreset weights  $w$  such that for each set of weights  $w$ ,  $\kappa_w$  has invariant distribution  $\pi_w$ . Then one step of Coreset MCMC, shown in Algorithm 1, involves the following updates in sequence:

$$w_{t+1} \leftarrow \text{proj}_{\mathcal{W}}(w_t - \gamma_t g(w_t, \theta_t, \mathcal{S}_t)) \quad (5)$$

$$\theta_{(t+1)k} \sim \kappa_{w_{t+1}}(\cdot | \theta_{tk}), \quad k = 1, \dots, K,$$

where  $\mathcal{S}_t$  is a subset of  $S$  indices from  $\{1, \dots, N\}$  drawn uniformly without replacement. Note that  $K \geq 2$  Markov chains are required for the gradient estimate in Eq. (3); each step in these  $K$  chains is independent given the current weights, and can be run in parallel.

**Initialization:** We initialize the coreset weights uniformly to have sum  $N$ :

$$w_{0m} = \frac{N}{M} \quad m = 1, \dots, M,$$

and initialize the Markov chain states  $\theta_{0k}$ ,  $k = 1, \dots, K$  arbitrarily, although it may help in practice to allow some amount of burn-in before beginning to adapt the coreset weights.

**Choice of kernel:** Coreset MCMC is agnostic as to the choice of kernels  $\kappa_w$ ; for example,  $\kappa_w$  could be a basic random-walk Metropolis–Hastings scheme (Robert and Casella, 2004, Ch. 7), slice sampler (Neal, 2003), or Hamiltonian Monte Carlo-based method (Neal, 2011; Duane et al., 1987) designed to target  $\pi_w$ . However, because the weights  $w_t$ —and hence the coreset posterior  $\pi_{w_t}$ —will change as iterations proceed, one should prefer kernels that automatically adapt to the local

geometry of the coreset posterior. For example, a hit-and-run slice sampler with doubling (Neal, 2003, Fig. 4; Bélisle et al., 1993) is a reasonable choice as it will automatically select reasonable step sizes as  $w$  changes. Although adaptive MCMC kernels (see Andrieu and Thoms, 2008) may seem attractive here, they are usually designed for a single target—not a moving target like  $\pi_{w_t}$ —and tuning such methods jointly with Coreset MCMC may be difficult.

**Complexity:** If the kernel  $\kappa_w$  can be applied in  $O(M)$  time, and each log-likelihood term  $\ell_n(\theta)$  takes  $O(1)$  time to compute, each step of Coreset MCMC takes  $O((M + S)K)$  time on a single processor. If  $K$  processors are available, this can be reduced to  $O((M + S) \log K)$  by running the Markov chains in parallel and using distributed reduction in Eq. (3).

**Subsampling:** We draw  $S_t$  without replacement, as this reduces subsampling variance by  $\frac{N-S}{N-1} \leq 1$  versus sampling with replacement. However, when  $S \ll N$ , the difference between these two is small. In practice, one can perform random order scans over the data.

### 3.3 Convergence Analysis

We now analyze the convergence behaviour of Coreset MCMC. Proofs may be found in Appendix A. So far we have assumed that  $\gamma_t > 0$  is a (not necessarily strictly) monotone decreasing sequence,  $\mathcal{W}$  is closed and convex, and that  $\forall w \in \mathcal{W}$ ,  $Z(w) < \infty$ ; all of these assumptions are required to use Algorithm 1. In this subsection, we impose a set of additional assumptions that are not required by Algorithm 1 but simplify the analysis.

Note that the coreset optimization problem Eq. (1) is nonconvex; but despite this fact, we are able to show that Coreset MCMC produces an optimal set of weights in settings where  $M$  is large enough such that there is an exact coreset. We formalize this assumption in Assumption 3.1.

**Assumption 3.1** (Exact coreset). *There exists a unique  $w^* \in \mathbb{R}^M$ ,  $c^* \in \mathbb{R}$  such that  $w^* \in \mathcal{W}$  and*

$$\sum_{n=1}^N \ell_n(\cdot) = \sum_{m=1}^M w_m^* \ell_m(\cdot) + c^* \quad \pi_0 - a.e.v.$$

We now begin with the simple case where  $S = N$ , i.e., where there is no data subsampling in the stochastic gradient estimate Eq. (3). Define  $G_t \in \mathbb{R}^{M \times M}$  to be

$$G_t = \frac{1}{K-1} \sum_{k=1}^K \begin{bmatrix} \bar{\ell}_1(\theta_{tk}) \\ \vdots \\ \bar{\ell}_M(\theta_{tk}) \end{bmatrix} \begin{bmatrix} \bar{\ell}_1(\theta_{tk}) \\ \vdots \\ \bar{\ell}_M(\theta_{tk}) \end{bmatrix}^\top,$$

and note that the KL gradient estimate in Eq. (3) can

be written

$$g(w_t, \theta_t, [N]) = G_t (w_t - w^*).$$

Hence all of the stochasticity in the gradient estimate in Eq. (3) comes from  $G_t$ , which estimates  $C_{w_t} = \text{Cov}_{\pi_{w_t}} [\ell_1(\theta) \dots \ell_M(\theta)]$ . As is usual in the analysis of stochastic optimization, we require two high-level conditions for convergence: the gradient estimate must (1) provide progress per iteration on average, and (2) not be so noisy that the algorithm makes unrecoverable mistakes. We formulate these in Assumptions 3.2 and 3.3 in terms of the moments of  $G_t$  after a single Markov chain step.

**Assumption 3.2** (Markov gradient mixing). *There exists  $\lambda > 0$  such that*

$$\forall w_t \in \mathcal{W}, \theta_{t-1} \in \Theta^K \quad \mathbb{E}[G_t \mid w_t, \theta_{t-1}] \succeq \lambda I.$$

**Assumption 3.3** (Markov gradient noise boundedness). *There exists a  $\bar{\lambda} < \infty$  such that*

$$\forall w_t \in \mathcal{W}, \theta_{t-1} \in \Theta^K \quad \mathbb{E}[G_t^\top G_t \mid w_t, \theta_{t-1}] \preceq \bar{\lambda} I.$$

Assumption 3.2 can be interpreted in two parts. First, we ask that the Markov chains mix quickly so that the covariance estimate  $G_t$  is similar to the exact covariance  $C_{w_t}$  in the KL gradient formula Eq. (2). Second, we require the exact covariance  $C_{w_t}$  to have a positive minimum eigenvalue to guarantee that Algorithm 1 progresses towards the optimum. Assumption 3.3 can be interpreted as placing a bound on the variance of  $G_t$  at each iteration. Both Assumptions 3.2 and 3.3 are meant to be representative, akin to the uniformly bounded gradient noise assumptions common in the optimization literature (e.g. Rakhlin et al., 2012; Hazan and Kale, 2014); they are too strong to hold precisely in realistic models with unbounded parameter spaces  $\Theta$ . However, note that Assumption 3.3 holds if the  $\ell_n$  are continuous and the parameter space  $\Theta$  is compact; and Assumption 3.2 holds if additionally  $\mathcal{W}$  is compact and the minimum eigenvalue of  $\mathbb{E}[G_t \mid w_t, \theta_{t-1}]$  is continuous and strictly positive for  $w_t \in \mathcal{W}$ ,  $\theta_{t-1} \in \Theta^K$ .

Theorem 3.4 shows that Coreset MCMC produces weights that converge to the optimum  $w^*$  in expectation (and hence in probability) when the full data are used for gradient estimation, i.e., when  $S = N$  in Eq. (3). In particular, a constant learning rate  $\gamma_t = \gamma$  in Theorem 3.4 yields linear convergence despite the use of stochastic gradient estimates. This is because the gradient variance shrinks as the iterates approach the optimum due to the multiplicative  $(w_t - w^*)$  factor in the gradient estimate. The link between variance reduction and linear convergence in stochastic gradient methods is well-known in the optimization literature; see Gower et al. (2020) for a more in-depth discussion.

**Theorem 3.4.** *Suppose Assumptions 3.1 to 3.3 hold and  $S = N$ . There exists  $\gamma > 0$  such that if  $\sup_t \gamma_t \leq \gamma$ ,*

$$\mathbb{E}\|w_t - w^*\|^2 \leq e^{-\lambda \sum_{\tau=0}^{t-1} \gamma_\tau} \|w_0 - w^*\|^2.$$

*In particular, if  $\sum_{t=0}^\infty \gamma_t = \infty$ , then  $\mathbb{E}\|w_t - w^*\|^2 \rightarrow 0$  as  $t \rightarrow \infty$ .*

When gradients are estimated using subsampling, i.e., when  $S < N$  in Eq. (3), there is an additional noise term that must be controlled in the analysis. Define

$$\begin{aligned} \Delta_{tkn} &= \bar{\ell}_n(\theta_{tk}) - \frac{1}{N} \sum_{n'=1}^N \bar{\ell}_{n'}(\theta_{tk}) \\ V_t &= \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \left( \frac{1}{K-1} \sum_{k=1}^K \bar{\ell}_m(\theta_{tk}) \Delta_{tkn} \right)^2. \end{aligned}$$

The quantity  $V_t \geq 0$  captures the variance of subsampling noise. Assumption 3.5 uniformly bounds the influence of this noise after a single Markov chain step. Again, although meant just to be representative, Assumption 3.5 holds at least when the log-likelihood functions are continuous and the parameter space  $\Theta$  is compact.

**Assumption 3.5** (Subsampling noise boundedness). *There exists a  $V < \infty$  such that*

$$\forall w_t \in \mathcal{W}, \theta_{t-1} \in \Theta^K \quad \mathbb{E}[V_t \mid w_t, \theta_{t-1}] \leq V.$$

Given the addition of Assumption 3.5, Theorem 3.6 provides a convergence guarantee for Coreset MCMC with data subsampling. Note that convergence is now sublinear due to subsampling noise.

**Theorem 3.6.** *Suppose Assumptions 3.1 to 3.3 and 3.5 hold. There exists  $\gamma > 0$  such that if  $\sup_t \gamma_t \leq \gamma$ , then*

$$\begin{aligned} \mathbb{E}\|w_t - w^*\|^2 &\leq e^{-\lambda \sum_{\tau=0}^{t-1} \gamma_\tau} \|w_0 - w^*\|^2 \\ &\quad + \frac{V(N-S)N^2M}{(N-1)S} \sum_{\tau=0}^{t-1} \gamma_\tau^2 e^{-\lambda \sum_{u=\tau+1}^{t-1} \gamma_u}. \end{aligned}$$

*In particular, if  $\sum_{t=0}^\infty \gamma_t = \infty$  and  $\gamma_t \rightarrow 0$  as  $t \rightarrow \infty$ , then  $\mathbb{E}\|w_t - w^*\|^2 \rightarrow 0$  as  $t \rightarrow \infty$ .*

### 3.4 Intuition from a Representative Model

In this subsection we build intuition on the behaviour of Algorithm 1 using a Gaussian location model in  $\mathbb{R}^d$ , with prior  $\theta \sim \mathcal{N}(0, I)$ , data  $X_n \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta, I)$ ,  $N = 10,000$ , and  $d = 20$ . We use the feasible region  $\mathcal{W} = \{w \in \mathbb{R}^M : w \geq 0, \sum_m w_m = N\}$ . The coreset posterior is  $\pi_w = \mathcal{N}(\mu_w, \sigma^2 I)$  with  $\mu_w =$

$\sigma^2 \sum_{m=1}^M w_m X_m$  and  $\sigma^2 = \frac{1}{1+N}$ . We use a family of Markov kernels  $\kappa_w$  inspired by Langevin dynamics,

$$\theta_{(t+1)k} \sim \mathcal{N}\left(\sqrt{\beta}(\theta_{tk} - \mu_{w_t}) + \mu_{w_t}, (1-\beta)I\right),$$

where  $\beta \in [0, 1]$  controls how quickly the Markov chains mix;  $\beta = 0$  provides independent draws,  $\beta = 1$  yields no mixing. We use a default setting of  $\beta = 0.8$ ,  $S = 30$ ,  $M = 30$ ,  $K = 20$  for each experiment, and vary one parameter while holding the others fixed. We use a learning rate of the form  $\gamma_t = \gamma(t+1)^{\alpha-1}$ , with  $\gamma = \frac{N}{10M}$  and—following Theorems 3.4 and 3.6— $\alpha = 1$  for full-data gradient estimates and  $\alpha = 0.5$  for subsampled estimates. We plot KL divergence computed using Eq. (10) versus “cost” equal to  $((M+S)\log K)t$ , per the earlier complexity analysis. A heuristic analysis of the expected KL divergence when  $\beta \approx 0$ ,  $N \gg M \gg 1$ , and  $\gamma \ll M$  (see Appendix B) yields the formula

$$\begin{aligned} \mathbb{E}D_{\text{KL}}(\pi_{w_t} \parallel \pi) &\lesssim \\ e^{-\frac{2\gamma(t^{\alpha-1})}{\alpha} \frac{N}{2M} + \frac{\gamma(N-S)d(K+d)(1+\log t^{1-\alpha})}{4S(K-1)t^{1-\alpha}}}. \end{aligned} \quad (6)$$

**Effect of parallelization ( $K$ ):** Inspecting the formula in Eq. (6), we expect increasing the number of chains should not influence the expected KL significantly for full-data gradients, and when subsampling should provide a meaningful reduction until  $K \approx d$ . Figs. 1a and 1b confirm this intuition. We recommend setting  $K$  as large as possible given available hardware in practice, as it may improve approximation quality without a noticeable time penalty.

**Effect of subsampling ( $S$ ):** Figs. 1a and 1b also demonstrate the influence of subsampled gradient estimates, and confirm the results in Eq. (6) and Theorems 3.4 and 3.6. In particular, the KL divergence appears to converge geometrically when using the full data gradient estimates, and polynomially when using subsampled estimates. However, adjusting for iteration cost, subsampling can provide improved approximation quality, especially in earlier iterations. A reasonable default setting of  $S$  in practice is such that the time spent simulating the Markov chains and computing the gradient update is comparable.

**Effect of coreset size ( $M$ ):** Fig. 1c shows that Coreset MCMC still converges when the optimal coreset is not exact, but that there is persisting error. Once the coreset gets large enough such that the optimal coreset is exact, further increasing the size does not yield major improvements. In this problem,  $M \approx d = 20$  suffices to get an exact coreset, and beyond that the effect of increasing  $M$  further diminishes.

**Effect of mixing rate ( $\beta$ ):** Fig. 1d shows that, at least in this example, the mixing rate of the Markov

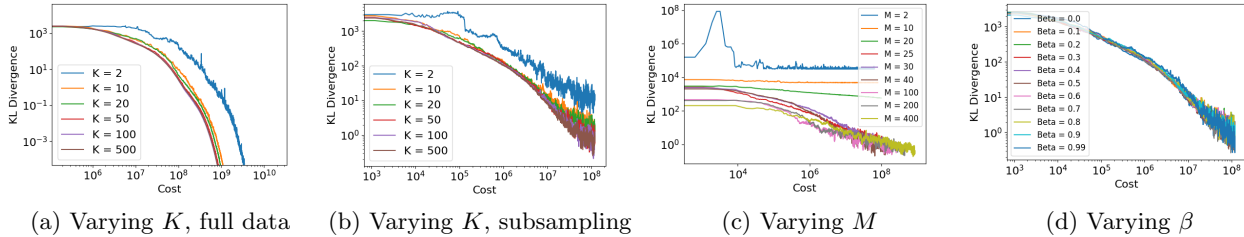


Figure 1: Influence of gradient estimation scheme (with or without subsampling), number of Markov chains, Markov chain mixing rate, and coreset size on the performance of Coreset MCMC in the Gaussian location model.

chain has little effect on the performance of Coreset MCMC. In practice, the mixing rate of any kernel  $\kappa_w$  could in principle be adjusted by increasing the number of steps between gradient updates. Given the results here, we recommend using a single application of an automatically adapting kernel, e.g., a slice sampler with doubling (Neal, 2003).

## 4 EXPERIMENTS

In this section, we compare **CoresetMCMC** (without subsampling, i.e.,  $S = N$ ) and **CoresetMCMC-S** (with subsampling, i.e.,  $S < N$ ) against previous coreset construction methods, as well as various subsampling MCMC methods. The coreset construction methods that we compare to are uniform subsampling (**Unif**)—which assigns a weight of  $N/M$  to  $M$  uniformly drawn points—as well as **QNC** (Naik et al., 2022) and **SHF** (Chen et al., 2022). Note that we do not compare to Sparse VI (Campbell and Beronov, 2019) due to its high computational cost; see Naik et al. (2022) for a comparison between Sparse VI and **QNC**. The subsampling MCMC methods we compare to are **Austerity MH** (**Austerity**) (Korattikara et al., 2014), **confidence MH** (**Confidence**) (Bardenet et al., 2014), and stochastic gradient Langevin dynamics (**SGLD-CV**) and stochastic gradient Hamiltonian Monte Carlo (**SGHMC-CV**) each with a control variate set to the log-likelihood gradient near the posterior density mode (Nemeth and Fearnhead, 2021).

We consider four Bayesian regression models: a (non-conjugate) linear, a logistic, a Poisson, and a sparse regression model. The models for the first three real data experiments contain only continuous variables, and that for the last synthetic experiment contains both continuous and discrete variables; see Appendix C for details. We use Stan (Carpenter et al., 2017) to obtain full data inference results for real data experiments, and use the Gibbs sampler developed by George and McCulloch (1993) for the synthetic experiment, due to its inclusion of discrete variables.

For the three real data experiments (with only continu-

ous variables), we measure the posterior approximation quality of all methods using the two-moment KL, defined as  $D_{\text{KL}}(\mathcal{N}(\hat{\mu}, \hat{\Sigma}) || \mathcal{N}(\mu, \Sigma))$ , where  $\hat{\mu}, \hat{\Sigma}$  are the mean and covariance estimated using draws from each method, and  $\mu, \Sigma$  are the same estimated for the full data posterior. This metric combines posterior mean and covariance error into a single number. We measure efficiency using wall-clock training time as well as the minimum marginal effective sample size (ESS) per second across all dimensions. For sparse regression (with both continuous and discrete variables), we measure the posterior approximation quality and sampling efficiency for the discrete and continuous posterior marginals separately. For the continuous variables, we use the same metrics as for the other experiments. For the discrete variables, we measure posterior approximation quality via the Jensen-Shannon divergence—which accounts for the possibility of differing support in empirical distribution approximations—and efficiency via the ESS computed for the fraction of correct feature inclusion indicators. For both ESS computations, we use the bulk-ESS formula in Vehtari et al. (2021).

For all coreset methods that use an MCMC kernel, we use the hit-and-run slice sampler with doubling (Bélisle et al., 1993; Neal, 2003) for linear and logistic regressions, the univariate slice sampler with doubling (Neal, 2003, Fig. 4) applied to each dimension for the more challenging Poisson regression, and the Gibbs sampler developed by George and McCulloch (1993) for sparse regression. For **SGLD-CV** and **SGHMC-CV**, we set the subsampling size to 500. For **Austerity** and **Confidence**, we set the accept-reject decision threshold to 0.05.

All experiments were performed on the UBC ARC Sockeye cluster. Each algorithm was run on 8 single-threaded cores of a 2.1GHz Intel Xeon Gold 6130 processor with 32GB memory. Code for these experiments is available at <https://github.com/NaitongChen/coreset-mcmc-experiments>. More experimental details and additional plots are in Appendices C and D.

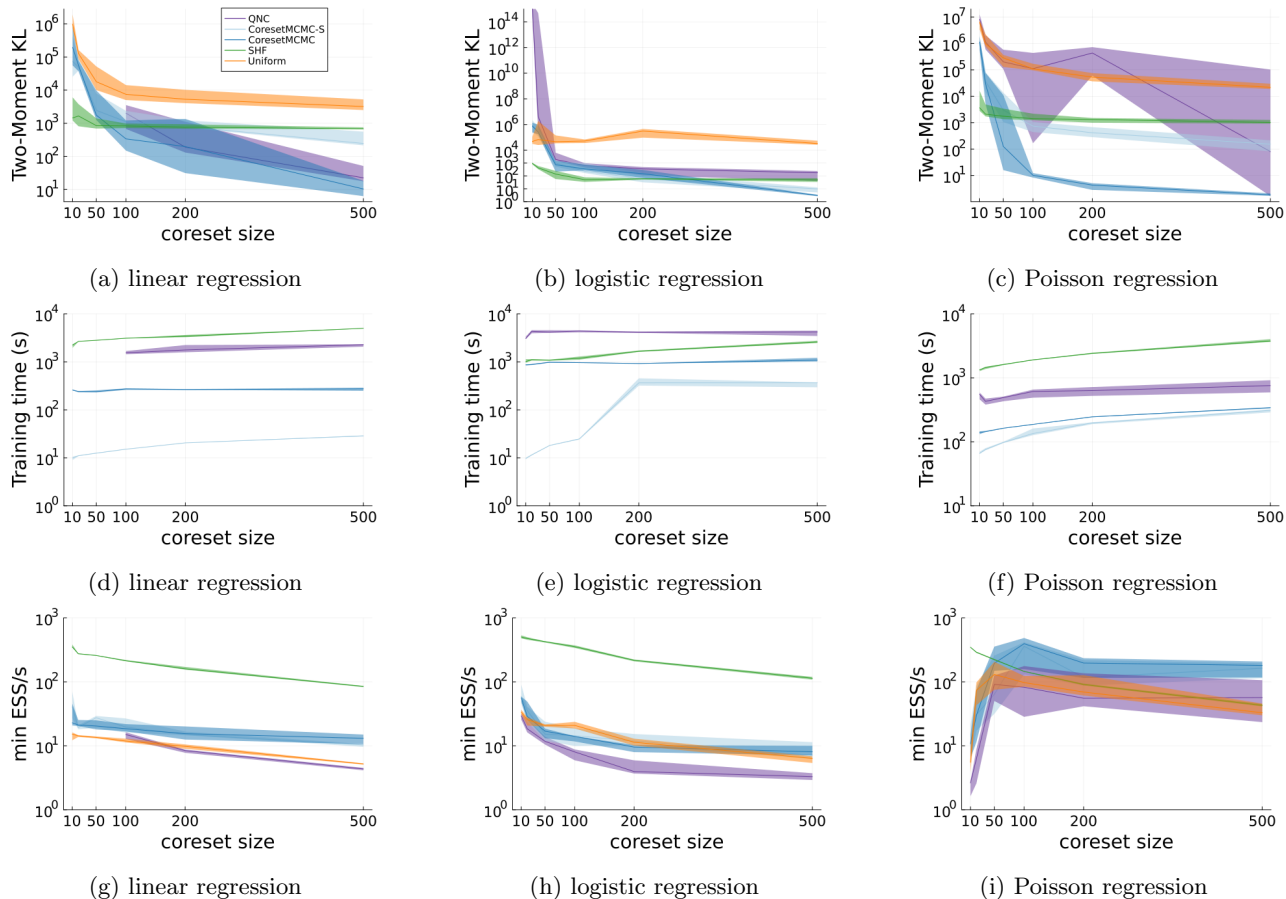


Figure 2: Comparison of coreset methods on real data examples. Figs. 2a to 2c show posterior approximation quality via the two-moment KL, Figs. 2d to 2f show training time, and Figs. 2g to 2i show sampling efficiency via the min. ESS per second. The lines indicate the median, and error regions indicate 25<sup>th</sup> to 75<sup>th</sup> percentile from 10 runs.

### 4.1 Comparison of Coreset Methods

Fig. 2 displays the comparison of Coreset MCMC with past coreset construction methods on the three real data experiments over various coreset sizes  $M$ . As shown in Figs. 2a to 2c, the `Unif` baseline, although not requiring any training and yielding a competitive sampling efficiency, generally provides poor quality coresets with an order of magnitude higher two-moment KL nearly uniformly across all coreset sizes than all other methods. `QNC` generates competitive posterior approximations for both linear and logistic regression, but a significantly worse approximation in the more challenging Poisson regression problem. Its approximations are also less reliable in general; the instability is particularly evident in linear regression when the coreset size is small—where `QNC` produces NaN values even after multiple attempts at tuning—and in the poisson regression example. `SHF` generally provides high quality approximations for small coreset sizes, but they do not improve with increasing coreset size. We conjecture

that this is due to the quasi-refreshment steps limiting the expressiveness of the variational family used in `SHF`.

Figs. 2d to 2f show that both `QNC` and `SHF` take between 2 to 10 times longer to train than both variants of Coreset MCMC. `QNC` takes longer than `CoresetMCMC` to train because it runs a full MCMC procedure at each optimization iteration, and computes a quasi-Newton update that has time complexity  $\mathcal{O}(M^3)$ ; compare to `CoresetMCMC`, which has updates that require only  $\mathcal{O}(M)$  time. `SHF` takes longer than `CoresetMCMC` to train because it requires taking the gradient over the entire flow to obtain updates on the coreset weights, refreshment parameters, and step sizes. Once trained, Figs. 2g to 2i show that all coreset methods provide a similar minimum ESS per second except for `SHF`, which is a variational methods that provides i.i.d. draws.

Fig. 3 shows the same comparison of coreset methods on sparse regression. Note that `SHF` is not applicable here due to the inclusion of discrete variables, and

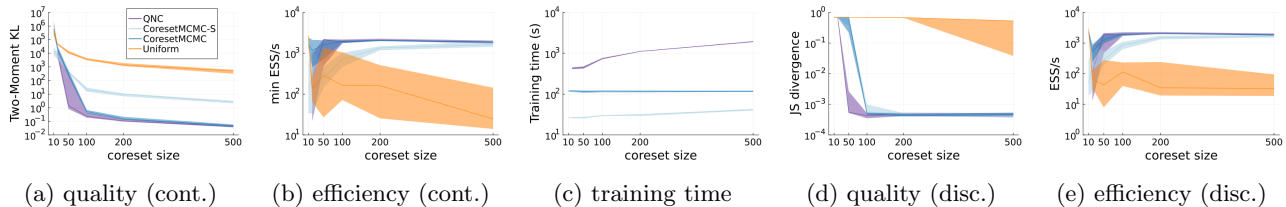


Figure 3: Comparison of coreset methods on sparse regression. Figs. 3a and 3b show posterior approximation quality and sampling efficiency across the continuous components via the two-moment KL and the min. ESS per second. Figs. 3d and 3e show those across the discrete components via the Jensen-Shannon divergence and ESS per second. Fig. 3c shows training time in seconds. The lines indicate the median, and error regions indicate 25<sup>th</sup> to 75<sup>th</sup> percentile from 10 runs.

hence is excluded from this experiment. We see that the trends in these plots resemble those from the three real data experiments: **Unif** produces worse quality coresets than **CoresetMCMC** and **QNC**, and **QNC** requires much longer training times than **CoresetMCMC**.

Comparing between the two variants of Coreset MCMC methods, it is clear that by including the full dataset when estimating the gradient, we obtain generally higher quality posterior approximations in terms of the two-moment KL. In general, we recommend setting  $S$  in Eq. (3) such that gradient updates require similar time as generating the next MCMC state, to balance time spent on simulating the Markov chains and estimating KL gradients.

## 4.2 Comparison with Subsampling MCMC

Figs. 4 and 5 display the comparison between all coreset methods and all subsampling MCMC methods, on the three real data experiments and the synthetic sparse regression experiment. Here we fix the coreset size to  $M = 500$ . On the easier linear and logistic regression problems (i.e., with more Gaussian-like posteriors), we see that **CoresetMCMC** provides approximations competitive, and sometimes better, than the other methods except for **SGHMC-CV**. However, we note that stochastic gradient MCMC methods depends heavily on the quality of the control variate. This is illustrated in the more challenging Poisson regression example; in this example, the control variate construction struggled to identify the mode of the posterior, causing resulting approximation quality to suffer. Furthermore, stochastic gradient MCMC methods are limited to models with only continuous variables, whereas **CoresetMCMC**, as illustrated in Fig. 5, applies to and performs well on models with both continuous and discrete variables.

In terms of sampling efficiency, **CoresetMCMC** and **CoresetMCMC-S** are uniformly better than other subsampling MCMC methods. Specifically, we get roughly 2 orders of magnitude higher ESS per second in **CoresetMCMC** than **SGHMC-CV**. This means that

**CoresetMCMC** can eventually catch up despite the initial training time. As an example, in the linear regression case, although it takes **CoresetMCMC** roughly 500 seconds to train, it is only enough time for **SGHMC-CV** to obtain roughly 50 effective samples. This deficit can be recovered by **CoresetMCMC** after 5 seconds once it stops adapting the coreset weights.

## 4.3 Tuning Difficulty

Although not reflected in the results in Figs. 2 to 5, tuning difficulty was a key differentiator in usability among the methods tested. Compared with other coreset methods, we found **CoresetMCMC** very straightforward to tune, as it just requires setting a learning rate. For **QNC**, in contrast, one needs to pick the number of MCMC samples used to estimate the gradient, the optimization step size, the number of iterations in which to perform a line search on the step size, and the number of weight optimization iterations. We note that the recommendation of setting the optimization step size to 1 from Naik et al. (2022) may cause the coreset weights to become unstable; we instead set this value to the step size at the last iteration where a line search was performed. Despite this modification, we can see from our results that **QNC** is still less stable than other coreset methods, and as such required significantly more manual tuning effort and time than **CoresetMCMC**. For **SHF**, much of the tuning effort goes into specifying the flow architecture: one must select the number of quasi-refreshment steps and the number of leapfrog steps in between quasi-refreshments. As discussed earlier, these tuning parameters determine the expressiveness of the variational family, and so one needs to balance having a highly expressive variational family without it being too complex to train within reasonable time.

For the stochastic gradient MCMC methods with control variates, the key difficulty (beyond tuning the subsample size) was to balance the effort spent constructing a good control variate (i.e., finding a reasonable estimate of the posterior mode) versus the effort spent



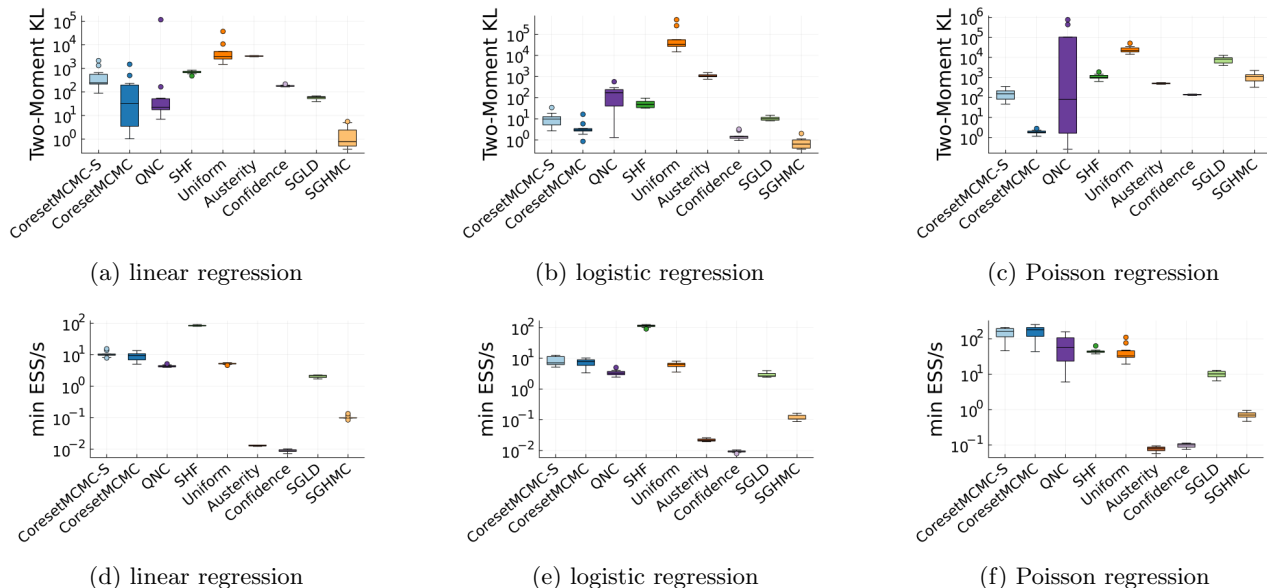


Figure 4: Comparison of coresets and subsampling MCMC methods on real data examples. Figs. 4a to 4c show posterior approximation quality via the two-moment KL, and Figs. 4d to 4f show sampling efficiency via the min. ESS per second. The boxplots indicate the median, 25<sup>th</sup>, and 75<sup>th</sup> percentiles from 10 runs.

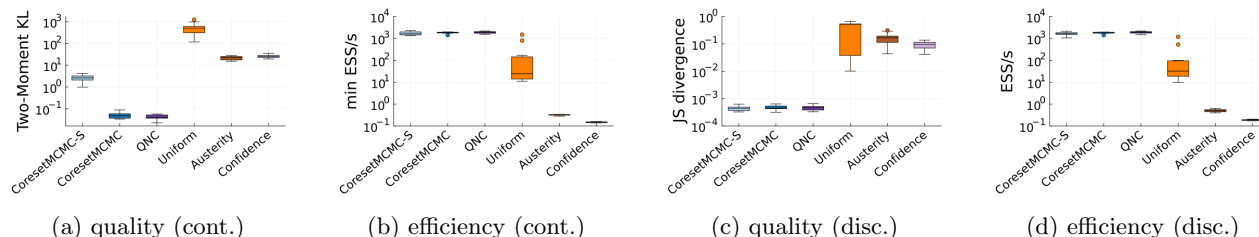


Figure 5: Comparison of coresets and subsampling MCMC methods on sparse regression. Figs. 5a and 5b show posterior approximation quality and sampling efficiency across the continuous components via the two-moment KL and the min. ESS per second. Figs. 5c and 5d show those across the discrete components via the Jensen-Shannon divergence and ESS per second. The boxplots indicate the median, 25<sup>th</sup>, and 75<sup>th</sup> percentiles from 10 runs.

using the control variate in sampling. For *SGHMC-CV*, one also needs to specify the number of leapfrog steps and the leapfrog step size. These two parameters are as complicated to tune as in Hamiltonian Monte Carlo.

Finally, both *Austerity* and *Confidence* require the user to input a threshold that determines the quality of the accept-reject decision approximation. This quantity does not directly influence the sampling efficiency as subsample size would. As a result, if one has a clear computation time budget, it may be more intricate to tune these methods well.

## 5 CONCLUSION

This paper introduced Coreset MCMC, a novel coreset construction method that is simple to implement and

tune compared with previous coreset methods. Coreset MCMC involves interleaving simulating states from several independent Markov chains targeting the coreset posterior with updating the coreset weights using these same draws. Theoretical results demonstrated that our method is asymptotically exact, assuming there exists a coreset whose corresponding posterior approximation has KL 0 to the true posterior, as well as other technical conditions. Finally, empirical results demonstrated that our method provides higher quality posterior approximations compared with other coreset methods, and improved sampling efficiency compared with subsampling MCMC methods.

## Acknowledgements

We acknowledge the support of an NSERC Discovery Grant (RGPIN-2019-03962), and the use of the ARC

Sockeye computing platform from the University of British Columbia.

## References

- Christian Robert and George Casella. *Monte Carlo Statistical Methods*. Springer, 2<sup>nd</sup> edition, 2004.
- Christian Robert and George Casella. A short history of Markov chain Monte Carlo: subjective recollections from incomplete data. *Statistical Science*, 26(1):102–115, 2011.
- Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin. *Bayesian Data Analysis*. CRC Press, 3<sup>rd</sup> edition, 2013.
- Marco Banterle, Clara Grazian, Anthony Lee, and Christian Robert. Accelerating Metropolis-Hastings algorithms by delayed acceptance. *Foundations of Data Science*, 1(2):103–128, 2019.
- Matias Quiroz, Robert Kohn, Mattias Villani, and Minh-Ngoc Tran. Speeding up MCMC by efficient data subsampling. *Journal of the American Statistical Association*, 114(526):831–843, 2019.
- Dougal Maclaurin and Ryan Adams. Firefly Monte Carlo: exact MCMC with subsets of data. In *Uncertainty in Artificial Intelligence*, 2014.
- Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in MCMC land: cutting the Metropolis-Hastings budget. In *International Conference on Machine Learning*, 2014.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *International Conference on Machine Learning*, 2014.
- Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine Learning*, 2011.
- Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International Conference on Machine Learning*, 2014.
- Jack Baker, Paul Fearnhead, Emily Fox, and Christopher Nemeth. `sgmcmc`: an R package for stochastic gradient Markov chain Monte Carlo. *Journal of Statistical Software*, 91(3):1–27, 2019.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. On Markov chain Monte Carlo methods for tall data. *Journal of Machine Learning Research*, 18(47):1–43, 2017.
- Matias Quiroz, Mattias Villani, Robert Kohn, Minh-Ngoc Tran, and Khue-Dung Dang. Subsampling MCMC - an introduction for the survey statistician. *Sankhya A*, 80(1):33–69, 2018.
- Christopher Nemeth and Paul Fearnhead. Stochastic gradient Markov chain Monte Carlo. *Journal of the American Statistical Association*, 116(533):433–450, 2021.
- Jonathan Huggins, Trevor Campbell, and Tamara Broderick. Coresets for scalable Bayesian logistic regression. In *Advances in Neural Information Processing Systems*, 2016.
- Trevor Campbell and Tamara Broderick. Automated scalable Bayesian inference via Hilbert coresets. *The Journal of Machine Learning Research*, 20(1):551–588, 2019.
- Trevor Campbell and Tamara Broderick. Bayesian coreset construction via greedy iterative geodesic ascent. In *International Conference on Machine Learning*, 2018.
- Trevor Campbell and Boyan Beronov. Sparse variational inference: Bayesian coresets from scratch. In *Advances in Neural Information Processing Systems*, 2019.
- Dionysis Manousakas, Zuheng Xu, Cecilia Mascolo, and Trevor Campbell. Bayesian pseudocoresets. In *Advances in Neural Information Processing Systems*, 2020.
- Cian Naik, Judith Rousseau, and Trevor Campbell. Fast Bayesian coresets via subsampling and quasi-Newton refinement. In *Advances in Neural Information Processing Systems*, 2022.
- Naitong Chen, Zuheng Xu, and Trevor Campbell. Bayesian inference via sparse Hamiltonian flows. In *Advances in Neural Information Processing Systems*, 2022.
- Jacky Zhang, Rajiv Khanna, Anastasios Kyrillidis, and Oluwasanmi Koyejo. Bayesian coresets: revisiting the nonconvex optimization perspective. In *Artificial Intelligence and Statistics*, 2021.
- Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer, 1<sup>st</sup> edition, 1990.
- Christophe Andrieu and Johannes Thoms. A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373, 2008.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- Léon Bottou. Stochastic learning. In Olivier Bousquet, Ulrike von Luxburg, and Gunnar Rätsch, editors, *Advanced Lectures on Machine Learning: ML Summer Schools 2003*, chapter 7, pages 146–168. Springer Berlin Heidelberg, 2004.

- Diederik Kingma and Jimmy Ba. Adam: a method for stochastic optimization. In *International Conference on Learning Representations*, 2014.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011.
- Radford Neal. Slice sampling. *The Annals of Statistics*, 31(3):705–767, 2003.
- Radford Neal. MCMC using Hamiltonian dynamics. In Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors, *Handbook of Markov chain Monte Carlo*, chapter 5. CRC Press, 2011.
- Simon Duane, Anthony Kennedy, Brian Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.
- Claude Bélisle, Edwin Romeijn, and Robert Smith. Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18(2):255–266, 1993.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making stochastic gradient descent optimal for strongly convex problems. In *International Conference on Machine Learning*, 2012.
- Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research*, 15(1):2489–2512, 2014.
- Robert Gower, Mark Schmidt, Francis Bach, and Peter Richtárik. Variance-reduced methods for machine learning. *Proceedings of the IEEE*, 108(11):1968–1983, 2020.
- Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: a probabilistic programming language. *Journal of Statistical Software*, 76(1):1–32, 2017.
- Edward George and Robert McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993.
- Aki Vehtari, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. Rank-normalization, folding, and localization: an improved  $\hat{R}$  for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2):667–718, 2021.
- Philip Wolfe. Convergence conditions for ascent methods. *SIAM Review*, 11(2):226–235, 1969.

## Checklist

1. For all models and algorithms presented, check if you include:

- (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. **Yes, this may be found in Section 3.**
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. **Yes, this may be found in Section 3, with more detailed in Appendices A and B.**
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. **Yes, this is included in the zip file in the supplementary material.**
2. For any theoretical claim, check if you include:
    - (a) Statements of the full set of assumptions of all theoretical results. **Yes, this may be found in Section 3.**
    - (b) Complete proofs of all theoretical results. **Yes, this may be found in Appendices A and B.**
    - (c) Clear explanations of any assumptions. **Yes, the assumptions are listed in Section 3, with more detailed explanations in Appendices A and B.**
  3. For all figures and tables that present empirical results, check if you include:
    - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). **Yes, this is included in the zip file in the supplementary material, and a URL is also provided in Section 4.**
    - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). **Yes, this may be found in Appendix C.**
    - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). **Yes, the error metrics are defined in Section 4, and the error bars are explained in the caption of each figure.**
    - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). **Yes, this may be found in Section 4.**
  4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
    - (a) Citations of the creator if your work uses existing assets. **Yes, these are included as footnotes in Appendix C.**

- (b) The license information of the assets, if applicable. **Not Applicable. We did not use any licensed assets.**
  - (c) New assets either in the supplemental material or as a URL, if applicable. **Yes, this is included in the zip file in the supplementary material, and a URL is also provided in Section 4.**
  - (d) Information about consent from data providers/curators. **Not Applicable. We did not use any data that required consent from providers/curators.**
  - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. **Not Applicable. None of such information is included in the paper.**
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. **Not Applicable. No human subjects are involved in this paper.**
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. **Not Applicable.**
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. **Not Applicable.**

## A PROOFS

*Proof of Theorem 3.6.* Define

$$H_t := \frac{1}{K-1} \sum_{k=1}^K \begin{bmatrix} \bar{\ell}_1(\theta_{tk}) \\ \vdots \\ \bar{\ell}_M(\theta_{tk}) \end{bmatrix} \begin{bmatrix} \bar{\ell}_1(\theta_{tk}) \\ \vdots \\ \bar{\ell}_N(\theta_{tk}) \end{bmatrix}^\top \in \mathbb{R}^{M \times N}.$$

Note that  $H_t$  is similar to  $G_t$  except that there are  $N$  columns. Define  $s_t \in \mathbb{R}^N$  such that  $\forall j \in \mathcal{S}_t, s_{tj} = \frac{N}{S}$  and 0 otherwise. The subsampled gradient estimate can then be written in the form

$$g(w_t, \theta_t, \mathcal{S}_t) = G_t(w_t - w^*) + H_t(1 - s_t).$$

We apply the projected gradient update to get

$$\begin{aligned} \|w_{t+1} - w^*\|^2 &= \|\text{proj}_{\mathcal{W}}(w_t - \gamma_t G_t(w_t - w^*) - \gamma_t H_t(1 - s_t)) - w^*\|^2 \\ &= \|\text{proj}_{\mathcal{W}}(w_t - \gamma_t G_t(w_t - w^*) - \gamma_t H_t(1 - s_t)) - \text{proj}_{\mathcal{W}} w^*\|^2 \\ &\leq \|w_t - \gamma_t G_t(w_t - w^*) - \gamma_t H_t(1 - s_t) - w^*\|^2 \\ &= \|(I - \gamma_t G_t)(w_t - w^*) - \gamma_t H_t(1 - s_t)\|^2. \end{aligned}$$

The second equality follows because  $w^* \in \mathcal{W}$  by assumption. The inequality follows because  $\mathcal{W}$  is convex and closed by assumption, and hence  $\text{proj}_{\mathcal{W}}$  is a contraction. Therefore,

$$\|w_{t+1} - w^*\|^2 \leq \|(I - \gamma_t G_t)(w_t - w^*) - \gamma_t H_t(1 - s_t)\|^2.$$

Expand the right hand side above and take the expectation, we get

$$\begin{aligned} \mathbb{E} \|w_{t+1} - w^*\|^2 &\leq \mathbb{E} [(w_t - w^*)^\top (I - \gamma_t G_t)^\top (I - \gamma_t G_t)(w_t - w^*)] + \gamma_t^2 \mathbb{E} [(1 - s_t)^\top H_t^\top H_t (1 - s_t)] \\ &\quad - 2\gamma_t \mathbb{E} [(w_t - w^*)^\top (I - \gamma_t G_t)^\top H_t (1 - s_t)]. \end{aligned} \tag{7}$$

Now use the tower property to show that the last term above is 0:

$$\begin{aligned} \mathbb{E} [(w_t - w^*)^\top (I - \gamma_t G_t)^\top H_t (1 - s_t)] &= \mathbb{E} [\mathbb{E} [(w_t - w^*)^\top (I - \gamma_t G_t)^\top H_t (1 - s_t) \mid w_t, \theta_t]] \\ &= \mathbb{E} [(w_t - w^*)^\top (I - \gamma_t G_t)^\top H_t \mathbb{E} [(1 - s_t) \mid w_t, \theta_t]] \\ &= 0, \end{aligned}$$

where the second line follows by noting  $\mathbb{E}[1 - s_t | w_t, \theta_t] = \mathbb{E}[1 - s_t] = 0$  due to the unbiased subsampling. We now focus on the second term on the right-hand side of Eq. (7). Again note that  $\mathbb{E}[s_t | \theta_t] = \mathbb{E}[s_t] = 1$ , so

$$\begin{aligned}
 & \mathbb{E}[(1 - s_t)^\top H_t^\top H_t(1 - s_t) | \theta_t] \\
 &= -1^\top H_t^\top H_t 1 + \mathbb{E}[s_t^\top H_t^\top H_t s_t | \theta_t] \\
 &= \mathbb{E}\left[\sum_{n, n'} s_{tn} s_{tn'} (H_t^\top H_t)_{nn'} | \theta_t\right] - \sum_{n, n'} (H_t^\top H_t)_{nn'} \\
 &= \frac{N}{S} \left( \sum_n \left(1 - \frac{S}{N}\right) (H_t^\top H_t)_{nn} + \sum_{n \neq n'} \left(\frac{S-1}{N-1} - \frac{S}{N}\right) (H_t^\top H_t)_{nn'} \right) \\
 &= \frac{N}{S} \left( \sum_n \left(\frac{N-S}{N}\right) (H_t^\top H_t)_{nn} - \sum_{n \neq n'} \frac{N-S}{N(N-1)} (H_t^\top H_t)_{nn'} \right) \\
 &= \frac{N(N-S)}{S(N-1)} \left( \sum_n (H_t^\top H_t)_{nn} - \frac{1}{N} \sum_{n, n'} (H_t^\top H_t)_{nn'} \right) \\
 &= \frac{N^2(N-S)}{S(N-1)(K-1)^2} \sum_{k, k'} \left( \sum_m \bar{\ell}_m(\theta_{tk}) \bar{\ell}_m(\theta_{tk'}) \right) \left( \frac{1}{N} \sum_n \bar{\ell}_n(\theta_{tk}) \bar{\ell}_n(\theta_{tk'}) - \frac{1}{N^2} \sum_{n, n'} \bar{\ell}_n(\theta_{tk}) \bar{\ell}_{n'}(\theta_{tk'}) \right) \\
 &= \frac{N^2(N-S)}{S(N-1)(K-1)^2} \sum_{k, k'} \left( \sum_m \bar{\ell}_m(\theta_{tk}) \bar{\ell}_m(\theta_{tk'}) \right) \left( \frac{1}{N} \sum_n \left( \bar{\ell}_n(\theta_{tk}) - \frac{1}{N} \sum_{n'} \bar{\ell}_{n'}(\theta_{tk}) \right) \left( \bar{\ell}_n(\theta_{tk'}) - \frac{1}{N} \sum_{n'} \bar{\ell}_{n'}(\theta_{tk'}) \right) \right) \\
 &= \frac{N^2 M(N-S)}{S(N-1)} V_t,
 \end{aligned}$$

where  $V_t$  is as defined in Assumption 3.5. Therefore,

$$\begin{aligned}
 \mathbb{E}[(1 - s_t)^\top H_t^\top H_t(1 - s_t)] &= \mathbb{E}[\mathbb{E}[(1 - s_t)^\top H_t^\top H_t(1 - s_t) | \theta_t]] \\
 &= \frac{N^2 M(N-S)}{S(N-1)} \mathbb{E}[V_t].
 \end{aligned}$$

Given the previous two steps, Eq. (7) becomes

$$\mathbb{E}\|w_{t+1} - w^*\|^2 \leq \mathbb{E}[(w_t - w^*)^\top (I - \gamma_t G_t)^\top (I - \gamma_t G_t)(w_t - w^*)] + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} \mathbb{E}[V_t].$$

By Assumptions 3.2 and 3.3, we can rewrite the first term above as

$$\begin{aligned}
 & \mathbb{E}[(w_t - w^*)^\top (I - \gamma_t G_t)^\top (I - \gamma_t G_t)(w_t - w^*)] \\
 &= \mathbb{E}[\mathbb{E}[(w_t - w^*)^\top (I - \gamma_t G_t)^\top (I - \gamma_t G_t)(w_t - w^*) | w_t, \theta_{t-1}]] \\
 &= \mathbb{E}[(w_t - w^*)^\top (I - 2\gamma_t \mathbb{E}[G_t | w_t, \theta_{t-1}] + \gamma_t^2 \mathbb{E}[G_t^\top G_t | w_t, \theta_{t-1}]) (w_t - w^*)] \\
 &\leq \mathbb{E}[(w_t - w^*)^\top (I - 2\gamma_t \lambda I + \gamma_t^2 \bar{\lambda} I) (w_t - w^*)].
 \end{aligned}$$

By Assumption 3.5, we have

$$\mathbb{E}[V_t] = \mathbb{E}[\mathbb{E}[V_t | w_t, \theta_{t-1}]] \leq V.$$

Therefore

$$\begin{aligned}
 \mathbb{E}\|w_{t+1} - w^*\|^2 &\leq \mathbb{E}[(w_t - w_t^*)^\top (I - 2\gamma_t \lambda I + \gamma_t^2 \bar{\lambda} I) (w_t - w_t^*)] + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} V \\
 &= \mathbb{E}[(1 - 2\gamma_t \lambda + \gamma_t^2 \bar{\lambda}) \|w_t - w^*\|^2] + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} V \\
 &= (1 - 2\gamma_t \lambda + \gamma_t^2 \bar{\lambda}) \mathbb{E}\|w_t - w^*\|^2 + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} V.
 \end{aligned}$$

Now let  $\gamma = \lambda/\bar{\lambda} > 0$ , and suppose  $\sup_t \gamma_t \leq \gamma$ . Then

$$\begin{aligned} \mathbb{E}\|w_{t+1} - w^*\|^2 &\leq (1 - \gamma_t \lambda) \mathbb{E}\|w_t - w^*\|^2 + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} V \\ &\leq e^{-\gamma_t \lambda} \mathbb{E}\|w_t - w^*\|^2 + \gamma_t^2 \frac{N^2 M(N-S)}{S(N-1)} V. \end{aligned}$$

Then expand the recursion to obtain the stated bound:

$$\mathbb{E}\|w_t - w^*\|^2 \leq e^{-\lambda \sum_{\tau=0}^{t-1} \gamma_\tau} \|w_0 - w^*\|^2 + \frac{N^2 M(N-S)V}{S(N-1)} \sum_{\tau=0}^{t-1} \gamma_\tau^2 e^{-\lambda \sum_{u=\tau+1}^{t-1} \gamma_u}. \quad (8)$$

Convergence of the first term to 0 as  $t \rightarrow \infty$  occurs as long as  $\sum_{t=0}^{\infty} \gamma_t = \infty$ . It remains to show that the second sum term converges to 0 as  $t \rightarrow \infty$  when additionally  $\gamma_t \rightarrow 0$ . Note that since  $\sum_{t=0}^{\infty} \gamma_t = \infty$ ,

$$\gamma_0^2 e^{-\lambda \sum_{u=1}^{t-1} \gamma_u} \rightarrow 0, \quad t \rightarrow \infty,$$

and so the 0-index term in the sum converges to 0. Therefore we can focus on the sum from indices 1 to  $t-1$ . Note that since  $\gamma_t$  is a monotone decreasing sequence, it can be expanded to a monotone decreasing continuous real function through linear interpolation, so

$$\sum_{u=\tau+1}^{t-1} \gamma_u \geq \int_{\tau+1}^t \gamma_u du,$$

and thus we have

$$\sum_{\tau=1}^{t-1} \gamma_\tau^2 e^{-\lambda \sum_{u=\tau+1}^t \gamma_u} \leq \sum_{\tau=1}^{t-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du}.$$

Next we split the sum into two sums separated by some arbitrary index  $1 < b < t$ ,

$$\sum_{\tau=1}^{t-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} = \sum_{\tau=1}^{b-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} + \sum_{\tau=b}^{t-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du}. \quad (9)$$

We now bound each of the above two terms. The second term in Eq. (9) above can be bounded by a sum with a monotone decreasing summand, which can then be bounded with an integral:

$$\sum_{\tau=b}^{t-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} \leq \sum_{\tau=b}^{t-1} \gamma_\tau^2 \leq \int_b^t \gamma_{\tau-1}^2 d\tau.$$

The first sum term in Eq. (9) has a summand that is a product of a monotone decreasing and a monotone increasing sequence. We can split this sum further by isolating the summand at  $\tau = 0$ :

$$\begin{aligned} \sum_{\tau=1}^{b-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} &= \sum_{\tau=1}^{b-1} \gamma_\tau^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} \\ &\leq \int_1^b \gamma_{\tau-1}^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} d\tau \\ &= e^{-\lambda \int_0^t \gamma_u du} \int_1^b \gamma_{\tau-1}^2 e^{\lambda \int_0^{\tau+1} \gamma_u du} d\tau \\ &= e^{-\lambda \int_0^t \gamma_u du} \int_1^b \gamma_{\tau-1}^2 e^{\lambda \int_0^{\tau-1} \gamma_u du} e^{\lambda \int_{\tau-1}^{\tau+1} \gamma_u du} d\tau \\ &\leq e^{-\lambda \int_0^t \gamma_u du} \int_1^b \gamma_{\tau-1}^2 e^{\lambda \int_0^{\tau-1} \gamma_u du} e^{\lambda \int_0^2 \gamma_u du} d\tau \\ &\leq e^{-\lambda \int_0^t \gamma_u du + \lambda 2\gamma} \int_1^b \gamma_{\tau-1}^2 e^{\lambda \int_0^{\tau-1} \gamma_u du} d\tau \\ &\leq \gamma e^{-\lambda \int_0^t \gamma_u du + \lambda 2\gamma} \int_1^b \gamma_{\tau-1} e^{\lambda \int_0^{\tau-1} \gamma_u du} d\tau, \end{aligned}$$

where all inequalities follow by noting that  $\gamma_t$  is monotone decreasing. Now apply the change of variables  $x = \int_0^{\tau-1} \gamma_u du$  with  $dx = \gamma_{\tau-1} d\tau$  to get

$$\int_1^b \gamma_{\tau-1} e^{\lambda \int_0^{\tau-1} \gamma_u du} d\tau = \int_0^{\int_0^{b-1} \gamma_u du} e^{\lambda x} dx = \frac{e^{\lambda \int_0^{b-1} \gamma_u du} - 1}{\lambda} \leq \lambda^{-1} e^{\lambda \int_0^b \gamma_u du}.$$

Therefore the first term in Eq. (9) becomes

$$\sum_{\tau=1}^{b-1} \gamma_{\tau}^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} \leq \gamma \lambda^{-1} e^{-\lambda \int_0^t \gamma_u du + \lambda 2\gamma + \lambda \int_0^b \gamma_u du} = \gamma \lambda^{-1} e^{2\lambda\gamma - \lambda \int_b^t \gamma_u du}.$$

Therefore we have that

$$\sum_{\tau=1}^{t-1} \gamma_{\tau}^2 e^{-\lambda \int_{\tau+1}^t \gamma_u du} \leq \int_b^t \gamma_{\tau-1}^2 d\tau + \gamma \lambda^{-1} e^{2\lambda\gamma - \lambda \int_b^t \gamma_u du}.$$

As long as we can write  $b$  as a monotone increasing function of  $t$  such that

$$\int_b^t \gamma_{\tau}^2 d\tau \rightarrow 0 \quad \int_b^t \gamma_{\tau} d\tau \rightarrow \infty,$$

we have the desired result that Eq. (8) converges to 0 as  $t \rightarrow \infty$ . Note that we can make  $\int_b^t \gamma_{\tau} d\tau$  diverge to infinity slower by making  $b$  increase faster in  $t$ ; we can achieve an arbitrarily slow rate of divergence with a quickly enough increasing  $b$ . Therefore since  $\gamma_t \rightarrow 0$  as  $t \rightarrow \infty$  by assumption, we can select  $b$  increasing slowly enough that  $\int_b^t \gamma_{\tau} d\tau \rightarrow \infty$ , but quickly enough such that

$$\int_b^t \gamma_{\tau}^2 d\tau \leq \gamma_b \int_b^t \gamma_{\tau} d\tau \rightarrow 0.$$

□

*Proof of Theorem 3.4.* This result can be viewed as a special case of Theorem 3.6 where  $N = S$ , and so the result can be obtained by following the steps of Theorem 3.6. Note that when  $N = S$ , we have that for all  $t$ ,  $H_t(1 - s_t) = 0$ . Therefore, we can obtain the desired result without Assumption 3.5 or requiring that  $\gamma_t \rightarrow 0$  as  $t \rightarrow \infty$ . □

## B GAUSSIAN LOCATION EXAMPLE

In this section we derive an approximate bound for  $\mathbb{E}D_{\text{KL}}(\pi_w || \pi)$  under the Gaussian location model described in Section 3.4 and using the update rule in Section 3.2. Here we use the feasible region

$$\mathcal{W} = \{w \in \mathbb{R}^M : w \geq 0, \sum_m w_m = N\}.$$

For better readability, we use a slightly different notation for the coreset. In particular, we use

$$Y = [Y_1 \quad \dots \quad Y_M] \in \mathbb{R}^{d \times M}$$

to denote the coreset points, which is a subset of the entire dataset

$$X = [X_1 \quad \dots \quad X_N] \in \mathbb{R}^{d \times N}.$$

The full data posterior, as well as the coreset posterior given weights  $w \in \mathbb{R}^M$  can then be written as

$$\pi = \mathcal{N}\left(\frac{1}{1+N} \sum_{n=1}^N X_n, \frac{1}{1+N} I\right), \quad \pi_w = \mathcal{N}\left(\frac{\sum_{m=1}^M w_m Y_m}{1 + \sum_{m=1}^M w_m}, \frac{1}{1 + \sum_{m=1}^M w_m} I\right).$$



Since both  $\pi$  and  $\pi_{w_t}$  are Gaussian, we have a closed-form expression for  $D_{\text{KL}}(\pi_w||\pi)$ :

$$\begin{aligned} D_{\text{KL}}(\pi_w||\pi) &= \frac{1}{2} \left( d \log \frac{1 + \sum_m w_m}{1 + N} - d + d \frac{1 + N}{1 + \sum_m w_m} + (1 + N) \left\| \frac{\sum_m w_m Y_m}{1 + \sum_m w_m} - \frac{\sum_n X_n}{1 + N} \right\|^2 \right) \\ &= \frac{1}{2} \left( d \log \frac{1 + \mathbf{1}^\top w}{1 + N} - d + d \frac{1 + N}{1 + \mathbf{1}^\top w} + \frac{1 + N}{(1 + \mathbf{1}^\top w)^2} \left\| Yw - \frac{1 + \mathbf{1}^\top w}{1 + N} X \mathbf{1} \right\|^2 \right) \\ &= \frac{1}{2(1 + N)} \|Yw - X \mathbf{1}\|^2, \end{aligned} \quad (10)$$

where the last line is by  $\mathbf{1}^\top w = N$ . To obtain an approximate bound for  $\mathbb{E}D_{\text{KL}}(\pi_w||\pi)$ , we look at the update rule of Coreset MCMC on  $w$ . By assumption,  $\beta \approx 0$ , and so we assume that draws of  $\theta$  obtained using the Markov kernel  $\kappa_w$  are approximately i.i.d. Then following Eq. (3), given  $\theta_1, \dots, \theta_K \stackrel{\text{iid}}{\sim} \pi_w$ , we can write the unbiased estimate of  $\nabla_w D_{\text{KL}}(\pi_w||\pi)$  as

$$\begin{aligned} g(w, \theta, [N]) &= \frac{-1}{K-1} \sum_{k=1}^K \left( \begin{bmatrix} \ell_1(\theta_k) \\ \vdots \\ \ell_M(\theta_k) \end{bmatrix} - \begin{bmatrix} \frac{1}{K} \sum_{j=1}^K \ell_1(\theta_j) \\ \vdots \\ \frac{1}{K} \sum_{j=1}^K \ell_M(\theta_j) \end{bmatrix} \right) \\ &\quad \left( \sum_{n=1}^N \ell_n(\theta_k) - \sum_{m=1}^M w_m \ell_m(\theta_k) - \frac{1}{K} \sum_{j=1}^K \left( \sum_{n=1}^N \ell_n(\theta_j) - \sum_{m=1}^M w_m \ell_m(\theta_j) \right) \right) \\ &= \frac{-1}{K-1} \sum_{k=1}^K \left( \begin{bmatrix} \ell_1(\theta_k) \\ \vdots \\ \ell_M(\theta_k) \end{bmatrix} - \begin{bmatrix} \frac{1}{K} \sum_{j=1}^K \ell_1(\theta_j) \\ \vdots \\ \frac{1}{K} \sum_{j=1}^K \ell_M(\theta_j) \end{bmatrix} \right) \\ &\quad \left( \sum_{n=1}^N \left( \ell_n(\theta_k) - \frac{1}{K} \sum_{j=1}^K \ell_n(\theta_j) \right) - \sum_{m=1}^M \left( w_m \ell_m(\theta_k) - \frac{1}{K} \sum_{j=1}^K w_m \ell_m(\theta_j) \right) \right), \end{aligned} \quad (11)$$

Under the Gaussian location model, we know that

$$\ell_n(\theta) = -\frac{1}{2} \|X_n - \theta\|^2.$$

We can then simplify Eq. (11). First note that

$$\begin{aligned} \left( \begin{bmatrix} \ell_1(\theta_k) \\ \vdots \\ \ell_M(\theta_k) \end{bmatrix} - \begin{bmatrix} \frac{1}{K} \sum_{j=1}^K \ell_1(\theta_j) \\ \vdots \\ \frac{1}{K} \sum_{j=1}^K \ell_M(\theta_j) \end{bmatrix} \right) &= -\frac{1}{2} \left( \begin{bmatrix} \|Y_1 - \theta_k\|^2 \\ \vdots \\ \|Y_M - \theta_k\|^2 \end{bmatrix} - \begin{bmatrix} \frac{1}{K} \sum_{j=1}^K \|Y_1 - \theta_j\|^2 \\ \vdots \\ \frac{1}{K} \sum_{j=1}^K \|Y_M - \theta_j\|^2 \end{bmatrix} \right) \\ &= -\frac{1}{2} \left( \begin{bmatrix} -2Y_1^\top \theta_k + \|\theta_k\|^2 \\ \vdots \\ -2Y_M^\top \theta_k + \|\theta_k\|^2 \end{bmatrix} - \begin{bmatrix} \frac{1}{K} \sum_{j=1}^K -2Y_1^\top \theta_j + \|\theta_j\|^2 \\ \vdots \\ \frac{1}{K} \sum_{j=1}^K -2Y_M^\top \theta_j + \|\theta_j\|^2 \end{bmatrix} \right) \\ &= -\frac{1}{2} \left( \begin{bmatrix} -2Y_1^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) + \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \\ \vdots \\ -2Y_M^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) + \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \end{bmatrix} \right) \\ &= Y^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) - \frac{1}{2} \left( \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \right) \mathbf{1}. \end{aligned}$$

We now look at the two summation terms in the last line of Eq. (11).

$$\begin{aligned} \sum_{n=1}^N \left( \ell_n(\theta_k) - \frac{1}{K} \sum_{j=1}^K \ell_n(\theta_j) \right) &= -\frac{1}{2} \sum_{n=1}^N \left( \|X_n - \theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|X_n - \theta_j\|^2 \right) \\ &= (X1)^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) - \frac{N}{2} \left( \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \right). \end{aligned}$$

Similarly,

$$\sum_{m=1}^M \left( w_m \ell_m(\theta_k) - \frac{1}{K} \sum_{j=1}^K w_m \ell_m(\theta_j) \right) = (Yw)^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) - \frac{1^\top w}{2} \left( \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \right).$$

Therefore

$$\begin{aligned} g(w, \theta, [N]) &= \frac{1}{K-1} \sum_{k=1}^K \left( Y^\top \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right) - \frac{1}{2} \left( \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \right) 1 \right) \\ &\quad \left( \left( \theta_k - \frac{1}{K} \sum_{j=1}^K \theta_j \right)^\top (Yw - X1) - \frac{1^\top w - N}{2} \left( \|\theta_k\|^2 - \frac{1}{K} \sum_{j=1}^K \|\theta_j\|^2 \right) \right). \end{aligned}$$

Since  $1^\top w = N$ , we can use  $P = (I - M^{-1}11^\top)$  to obtain the projected gradient update as in Section 3.2 to get

$$\begin{aligned} w_{t+1} &= w_t - \gamma_t P g(w_t, \theta_t, [N]) \\ &= w_t - \gamma_t P Y^\top \frac{1}{K-1} \sum_{k=1}^K \left( \theta_{tk} - \frac{1}{K} \sum_{j=1}^K \theta_{tj} \right) \left( \theta_{tk} - \frac{1}{K} \sum_{j=1}^K \theta_{tj} \right)^\top (Yw_t - X1) \\ &= w_t - \gamma_t P Y^\top Q_t (Yw_t - X1), \end{aligned}$$

where

$$Q_t := \frac{1}{K-1} \sum_{k=1}^K \left( \theta_{tk} - \frac{1}{K} \sum_{j=1}^K \theta_{tj} \right) \left( \theta_{tk} - \frac{1}{K} \sum_{j=1}^K \theta_{tj} \right)^\top.$$

Note that the analysis so far assumes that one does not subsample the data to estimate the gradient. If one were to replace the full data in  $g$  with an unbiased subsample of the data, the projected gradient update of Coreset MCMC becomes

$$w_{t+1} = w_t - \gamma_t P Y^\top Q_t (Yw_t - S_t), \quad (12)$$

where we define  $S_t$  to be some random vector, independent of all else, such that  $\mathbb{E}S_t = X1$  and  $\text{Cov} S_t = \bar{\Sigma}$ . By further defining  $Z_t := Yw_t - X1$ , and  $A := P Y^\top$ , the recursion in Eq. (12) can instead be written in terms of  $Z_t$ :

$$Z_{t+1} = (I - \gamma_t A Q_t) Z_t + \gamma_t A Q_t (S_t - X1).$$

Solving the recursion, we have that

$$Z_t = \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right] Z_0 + \sum_{\tau=0}^{t-1} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] \gamma_\tau A Q_\tau (S_\tau - X1),$$

where matrix products indicate left multiplication with increasing index. We can also rewrite Eq. (10) as

$$\mathbb{E}D_{\text{KL}}(\pi_{w_t} \|\pi) = \frac{1}{2(1+N)} \mathbb{E} \|Z_t\|^2 = \frac{1}{2(1+N)} \text{tr} \mathbb{E} Z_t Z_t^\top. \quad (13)$$

The rest of the analysis approximates  $\mathbb{E}Z_t Z_t^\top$  to obtain a final approximate bound for  $\mathbb{E}D_{\text{KL}}(\pi_{w_t}||\pi)$ . With  $\boldsymbol{\theta}_t$  denoting the set of  $K$  samples at iteration  $t$ , we note that

$$\mathbb{E}[S_t - X1 \mid \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{t-1}] = \mathbb{E}[S_t - X1] = 0.$$

At the same time,  $S_t$  is independent across all  $t$  iterations. Therefore, we can expand  $\mathbb{E}Z_t Z_t^\top$  to get

$$\begin{aligned} \mathbb{E}Z_t Z_t^\top &= \mathbb{E} \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right] Z_0 Z_0^\top \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right]^\top \\ &\quad + \mathbb{E} \sum_{\tau=0}^{t-1} \gamma_\tau^2 \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau (S_\tau - X1)(S_\tau - X1)^\top Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &= \mathbb{E} \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right] Z_0 Z_0^\top \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right]^\top \\ &\quad + \sum_{\tau=0}^{t-1} \gamma_\tau^2 \mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau \bar{\Sigma} Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top. \end{aligned} \quad (14)$$

We now use the tower property to rewrite the expectation in the second term in Eq. (14) above to get

$$\begin{aligned} &\mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau \bar{\Sigma} Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &= \mathbb{E} \left[ \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A \mathbb{E} [Q_\tau \bar{\Sigma} Q_\tau^\top \mid \boldsymbol{\theta}_{\tau+1}, \dots, \boldsymbol{\theta}_{t-1}] A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \right]. \end{aligned}$$

Looking at the inner conditional expectation, we see that all randomness comes from  $\boldsymbol{\theta}_\tau$ . By letting

$$\bar{\boldsymbol{\theta}}_\tau = \frac{1}{K} \sum_{j=1}^K \boldsymbol{\theta}_{\tau j},$$

we then have

$$\begin{aligned} &\mathbb{E}_{\boldsymbol{\theta}_\tau} Q_\tau^\top \bar{\Sigma} Q_\tau \\ &= \frac{1}{(1+N)^2} \mathbb{E}_{\boldsymbol{\theta}_\tau} \frac{K^2}{(K-1)^2} \left( \frac{1}{K} \sum_{k=1}^K \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k}^\top - \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \right) \bar{\Sigma} \left( \frac{1}{K} \sum_{k=1}^K \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k}^\top - \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \right) \\ &= \frac{K^2}{(K-1)^2 (1+N)^2} \mathbb{E}_{\boldsymbol{\theta}_\tau} \left( \frac{1}{K^2} \sum_{k, k'=1}^K \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k}^\top \bar{\Sigma} \boldsymbol{\theta}_{\tau k'} \boldsymbol{\theta}_{\tau k'}^\top - \frac{1}{K} \sum_{k=1}^K \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k}^\top \bar{\Sigma} \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top - \frac{1}{K} \sum_{k=1}^K \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \bar{\Sigma} \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k}^\top + \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \bar{\Sigma} \bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \right) \\ &= \frac{K \bar{\Sigma} + \text{tr}(\bar{\Sigma}) I}{(K-1)(1+N)^2}, \end{aligned}$$

where the last equality is obtained by noting that  $\boldsymbol{\theta}_{\tau 1}, \dots, \boldsymbol{\theta}_{\tau k}$  are i.i.d. isotropic Gaussian with variance  $1/(1+N)$  and that

$$\mathbb{E} [\bar{\boldsymbol{\theta}}_\tau \bar{\boldsymbol{\theta}}_\tau^\top \mid \boldsymbol{\theta}_{\tau j}] = \mathbb{E} \left[ \frac{1}{K^2} \sum_{k, k'=1}^K \boldsymbol{\theta}_{\tau k} \boldsymbol{\theta}_{\tau k'}^\top \mid \boldsymbol{\theta}_{\tau j} \right] = \frac{1}{K^2} \boldsymbol{\theta}_{\tau j} \boldsymbol{\theta}_{\tau j}^\top + \frac{K-1}{K^2} I.$$

Therefore,

$$\begin{aligned} &\mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau \bar{\Sigma} Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &= \mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A \frac{K \bar{\Sigma} + \text{tr}(\bar{\Sigma}) I}{(K-1)(1+N)^2} A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top. \end{aligned}$$

Define  $B := A \frac{K\bar{\Sigma} + \text{tr} \bar{\Sigma} I}{(K-1)(1+N)^2} A^\top$ . Note that  $B$  is constant, and so we can write

$$\begin{aligned} \mathbb{E}(I - \gamma A Q_t) B (I - \gamma A Q_t)^\top &= \mathbb{E} \left[ B - \gamma A Q_t B - \gamma B Q_t^\top A^\top + \gamma^2 A Q_t B Q_t^\top A^\top \right] \\ &= B - \frac{\gamma}{1+N} AB - \frac{\gamma}{1+N} BA^\top + \gamma^2 A \mathbb{E} [Q_t B Q_t^\top] A^\top \\ &= B - \frac{\gamma}{1+N} AB - \frac{\gamma}{1+N} BA^\top + \gamma^2 A \left( \frac{KB + \text{tr}(B)I}{(K-1)(1+N)^2} \right) A^\top \\ &= \left( I - \frac{\gamma A}{1+N} \right) B \left( I - \frac{\gamma A}{1+N} \right)^\top + \gamma^2 \frac{ABA^\top + \text{tr}(B)AA^\top}{(K-1)(1+N)^2}, \end{aligned}$$

where the second equality is by noting that  $\mathbb{E}Q_t = \frac{1}{N+1}I$ . Since  $A$  is symmetric and positive semidefinite, we can write  $A = UDU^\top$  with  $D \succeq 0$  diagonal and  $U$  unitary. Then by defining  $E := U^\top A \frac{K\bar{\Sigma} + \text{tr} \bar{\Sigma} I}{(K-1)(1+N)^2} A^\top U$ , we have

$$B = UU^\top A \frac{K\bar{\Sigma} + \text{tr} \bar{\Sigma} I}{(K-1)(1+N)^2} A^\top UU^\top = UEU^\top.$$

Therefore,

$$\mathbb{E}(I - \gamma A Q_t) B (I - \gamma A Q_t)^\top = U \left( \left( I - \frac{\gamma}{1+N} D \right) E \left( I - \frac{\gamma}{1+N} D \right) + \gamma^2 \frac{DED + \text{tr}(E)D^2}{(K-1)(1+N)^2} \right) U^\top.$$

From this, we can see that starting with  $B = UEU^\top$ , the above operation yields

$$\mathbb{E}(I - \gamma A Q_t) B (I - \gamma A Q_t)^\top = \mathbb{E}(I - \gamma A Q_t) UEU^\top (I - \gamma A Q_t)^\top = UE'U^\top,$$

where each element of  $E'$  can be written in terms of  $E$  as

$$E'_{ij} = \left\{ \left( 1 - \frac{\gamma}{N+1} D_i \right) \left( 1 - \frac{\gamma}{N+1} D_j \right) + \frac{\gamma^2}{(K-1)(N+1)^2} D_i D_j \right\} E_{ij} + \frac{\gamma^2}{(K-1)(N+1)^2} \mathbb{1}[i=j] D_i^2 \text{tr} E.$$

Therefore, by applying this operation multiple times, we get

$$\begin{aligned} &\mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau \bar{\Sigma} Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &= \mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] UEU^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &= UE''U^\top, \end{aligned}$$

where

$$\begin{aligned} E''_{ij} &= \left[ \prod_{u=\tau+1}^{t-1} \left\{ \left( 1 - \frac{\gamma_u}{N+1} D_i \right) \left( 1 - \frac{\gamma_u}{N+1} D_j \right) + \frac{\gamma_u^2}{(K-1)(N+1)^2} D_i D_j \right\} \right] E_{ij} \\ &\quad + \left[ \prod_{u=\tau+1}^{t-1} \frac{\gamma_u^2}{(K-1)(N+1)^2} \mathbb{1}[i=j] D_i^2 \right] \text{tr} E. \end{aligned}$$

Considering  $\gamma_u$  small enough such that  $\gamma_u^2/(N+1)^2 \ll \gamma_u/(N+1)$ , we drop the higher order terms to find that

$$\begin{aligned} &\mathbb{E} \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right] A Q_\tau \bar{\Sigma} Q_\tau^\top A^\top \left[ \prod_{u=\tau+1}^{t-1} (I - \gamma_u A Q_u) \right]^\top \\ &\approx U \left[ \prod_{u=\tau+1}^{t-1} \left( I - \frac{\gamma_u}{N+1} D \right) \right] U^\top A \frac{K\bar{\Sigma} + \text{tr} \bar{\Sigma} I}{(K-1)(1+N)^2} A^\top U \left[ \prod_{u=\tau+1}^{t-1} \left( I - \frac{\gamma_u}{N+1} D \right) \right]^\top U^\top \\ &\approx e^{-\frac{\sum_{u=\tau+1}^{t-1} \gamma_u}{N+1}} A \frac{K\bar{\Sigma} + \text{tr} \bar{\Sigma} I}{(K-1)(1+N)^2} A^\top e^{-\frac{\sum_{u=\tau+1}^{t-1} \gamma_u}{N+1}} A^\top. \end{aligned}$$

For the expectation in the first term in Eq. (14), since  $Z_0 Z_0^\top$  is constant, we can use the same trick to get

$$\mathbb{E} \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right] Z_0 Z_0^\top \left[ \prod_{\tau=0}^{t-1} (I - \gamma_\tau A Q_\tau) \right]^\top \approx e^{-\frac{\sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} A} Z_0 Z_0^\top e^{-\frac{\sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} A^\top}.$$

Therefore,

$$\mathbb{E} Z_t Z_t^\top \approx e^{-\frac{\sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} A} Z_0 Z_0^\top e^{-\frac{\sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} A^\top} + \sum_{\tau=0}^{t-1} \frac{\gamma_\tau^2 e^{-\frac{\sum_{u=\tau+1}^{t-1} \gamma_u}{N+1} A} A (K \bar{\Sigma} + \text{tr} \bar{\Sigma} I) A^\top e^{-\frac{\sum_{u=\tau+1}^{t-1} \gamma_u}{N+1} A^\top}}{(K-1)(N+1)^2}.$$

Note that for uniform independent subsamples of size  $S$  without replacement,

$$\bar{\Sigma} = \frac{N^2(N-S)}{S(N-1)} \left[ \frac{1}{N} X X^\top - \frac{1}{N^2} X 1 1^\top X^\top \right] = \frac{N^2(N-S)}{S(N-1)} \Sigma_X,$$

where  $\Sigma_X = \frac{1}{N} X X^\top - \frac{1}{N^2} X 1 1^\top X^\top = \frac{1}{N} X (I - \frac{1}{N} 1 1^\top) X^\top$  is the empirical covariance of the full data. Similarly we have  $A = M \Sigma_Y$  for the empirical covariance of the coresets. Therefore,

$$\begin{aligned} \mathbb{E} Z_t Z_t^\top &\approx e^{-\frac{M \sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} \Sigma_Y} Z_0 Z_0^\top e^{-\frac{M \sum_{\tau=0}^{t-1} \gamma_\tau}{N+1} \Sigma_Y^\top} + \\ &\sum_{\tau=0}^{s-1} \frac{\gamma_\tau^2 M^2 N^2 (N-S) e^{-\frac{M \sum_{u=\tau+1}^{t-1} \gamma_u}{N+1} \Sigma_Y} \Sigma_Y (K \Sigma_X + \text{tr} \Sigma_X I) \Sigma_Y^\top e^{-\frac{M \sum_{u=\tau+1}^{t-1} \gamma_u}{N+1} \Sigma_Y^\top}}{(N-1)S(K-1)(N+1)^2}. \end{aligned}$$

By assumption,  $N \gg M \gg 1$ . Since the data are generated i.i.d. from a standard Gaussian, we can approximate  $\Sigma_X \approx \Sigma_Y \approx I$ . Therefore,

$$\mathbb{E} Z_t Z_t^\top \approx e^{-\frac{M(\sum_{\tau=0}^{t-1} \gamma_\tau + \sum_{\tau=0}^{t-1} \gamma_\tau)}{N+1}} Z_0 Z_0^\top + \frac{M^2(N-S)N^2(K+d)}{(N-1)S(K-1)(N+1)^2} \sum_{\tau=0}^{t-1} \gamma_\tau^2 e^{-\frac{M(\sum_{u=\tau+1}^{t-1} \gamma_u + \sum_{u=\tau+1}^{t-1} \gamma_u)}{N+1}} I.$$

We assume that  $\gamma_t = \gamma(t+1)^{\alpha-1}$  for some  $\gamma > 0$  and  $0 \leq \alpha \leq 1$ . Then we have the approximation

$$\sum_{\tau=0}^{t-1} \gamma_\tau \approx \gamma \int_0^{t-1} (\tau+1)^{\alpha-1} d\tau = \frac{\gamma(t^\alpha - 1)}{\alpha}.$$

Using this approximation and the fact that  $N+1 \approx N-1 \approx N$ , we obtain

$$\begin{aligned} \mathbb{E} Z_t Z_t^\top &\approx e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2})}{\alpha N}} Z_0 Z_0^\top + \frac{(N-S)M^2(K+d)}{NS(K-1)} \sum_{\tau=1}^t \gamma_{\tau-1}^2 e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2} - 2\tau^\alpha)}{\alpha N}} I \\ &= e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2})}{\alpha N}} Z_0 Z_0^\top + \frac{\gamma^2(N-S)M^2(K+d)}{NS(K-1)} \sum_{\tau=1}^t \tau^{2(\alpha-1)} e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2} - 2\tau^\alpha)}{\alpha N}} I. \end{aligned}$$

Again using an approximation of the sum as an integral, we get

$$\mathbb{E} Z_t Z_t^\top \approx e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2})}{\alpha N}} Z_0 Z_0^\top + \frac{\gamma^2(N-S)M^2(K+d)}{NS(K-1)} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2} - 2\tau^\alpha)}{\alpha N}} d\tau I.$$

We can now substitute the above expression to Eq. (13) to get

$$\begin{aligned} \mathbb{E} D_{\text{KL}}(\pi_{w_t} || \pi) &= \frac{1}{2(N+1)} \text{tr} \mathbb{E} Z_t Z_t^\top \\ &\approx \frac{1}{2(N+1)} \mathbb{E} \|Z_0\|^2 + \frac{\gamma^2(N-S)M^2(K+d)}{NS(K-1)} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2} - 2\tau^\alpha)}{\alpha N}} d\tau \text{tr} I \\ &\approx \frac{1}{2(N+1)} \mathbb{E} \|Z_0\|^2 + \frac{\gamma^2(N-S)M^2(K+d)d}{NS(K-1)} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{\gamma M(t^\alpha + t^{\alpha-2} - 2\tau^\alpha)}{\alpha N}} d\tau. \end{aligned}$$

Again by the assumption that  $N \gg M \gg 1$  and that all data points are i.i.d. standard Gaussian, we have  $\mathbb{E}\|Z_0\|^2 \approx N^2/M$ . Therefore,

$$\mathbb{E}D_{\text{KL}}(\pi_{w_t}|\pi) \approx e^{-\frac{2\gamma M(t^\alpha-1)}{\alpha N}} \frac{N}{2M} + \frac{\gamma^2(N-S)M^2(K+d)d}{2N^2S(K-1)} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{2\gamma M(t^\alpha-\tau^\alpha)}{\alpha N}} d\tau.$$

Set  $\gamma = cN/M$  for some tunable  $0 < c \ll M$ , then

$$\mathbb{E}D_{\text{KL}}(\pi_{w_t}|\pi) \approx e^{-\frac{2c(t^\alpha-1)}{\alpha}} \frac{N}{2M} + \frac{c^2(N-S)(K+d)d}{2S(K-1)} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{2c(t^\alpha-\tau^\alpha)}{\alpha}} d\tau. \quad (15)$$

We now derive an approximate upper bound for the integral term in Eq. (15) above. Apply the transformation of variables  $x = \tau^\alpha$  with  $dx = \alpha\tau^{\alpha-1}d\tau$  to get

$$\begin{aligned} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{2c(t^\alpha-\tau^\alpha)}{\alpha}} d\tau &= e^{-\frac{2ct^\alpha}{\alpha}} \int_1^t \tau^{2(\alpha-1)} e^{\frac{2c\tau^\alpha}{\alpha}} d\tau \\ &= \alpha^{-1} e^{-\frac{2ct^\alpha}{\alpha}} \int_1^{t^\alpha} x^{1-\alpha^{-1}} e^{\frac{2cx}{\alpha}} dx. \end{aligned}$$

Let  $1 \leq b \leq t$  and  $b \rightarrow \infty$  as  $t \rightarrow \infty$ , and so

$$\int_1^{b^\alpha} x^{1-\alpha^{-1}} e^{\frac{2cx}{\alpha}} dx \leq \int_1^{b^\alpha} e^{\frac{2cx}{\alpha}} dx \leq \frac{\alpha}{2c} \left( e^{\frac{2cb^\alpha}{\alpha}} - e^{\frac{2c}{\alpha}} \right) \leq \frac{\alpha}{2c} e^{\frac{2cb^\alpha}{\alpha}}.$$

We also have that

$$\int_{b^\alpha}^{t^\alpha} x^{1-\alpha^{-1}} e^{\frac{2cx}{\alpha}} dx \leq e^{\frac{2ct^\alpha}{\alpha}} \int_{b^\alpha}^{t^\alpha} x^{1-\alpha^{-1}} dx = e^{\frac{2ct^\alpha}{\alpha}} \frac{t^{2\alpha-1} - b^{2\alpha-1}}{2 - \alpha^{-1}}.$$

Therefore for any  $1 \leq b \leq t$ ,

$$\begin{aligned} \int_1^t \tau^{2(\alpha-1)} e^{-\frac{2c(t^\alpha-\tau^\alpha)}{\alpha}} d\tau &\leq \alpha^{-1} e^{-\frac{2ct^\alpha}{\alpha}} \left( \frac{\alpha}{2c} e^{\frac{2cb^\alpha}{\alpha}} + e^{\frac{2ct^\alpha}{\alpha}} \frac{t^{2\alpha-1} - b^{2\alpha-1}}{2 - \alpha^{-1}} \right) \\ &= \frac{1}{2c} e^{-\frac{2c(t^\alpha-b^\alpha)}{\alpha}} + \frac{t^{2\alpha-1} - b^{2\alpha-1}}{2\alpha - 1}. \end{aligned}$$

Then let  $b = t - h(t)$  where  $h(t) = o(t)$ . So as  $t$  gets large,  $b \approx t$ , and so

$$\begin{aligned} \frac{1}{2c} e^{-\frac{2c(t^\alpha-b^\alpha)}{\alpha}} + \frac{t^{2\alpha-1} - b^{2\alpha-1}}{2\alpha - 1} &\approx \frac{1}{2c} e^{-\frac{2ct^{\alpha-1}h(t)}{\alpha}} + \frac{t^{2\alpha-2}h(t)}{2\alpha - 1} \\ &\leq \frac{1}{2c} e^{-2ct^{\alpha-1}h(t)} + t^{2\alpha-2}h(t). \end{aligned}$$

By setting  $h(t) = \frac{1-\alpha}{2c} t^{1-\alpha} \log t$ ,

$$\int_1^t \tau^{2(\alpha-1)} e^{-\frac{2c(t^\alpha-\tau^\alpha)}{\alpha}} d\tau \lesssim \frac{1}{2ct^{1-\alpha}} + \frac{1-\alpha}{2ct^{1-\alpha}} \log t = \frac{1 + (1-\alpha) \log t}{2ct^{1-\alpha}}.$$

By substituting the above approximate bound to Eq. (15), we arrive at the final approximate bound

$$\mathbb{E}D_{\text{KL}}(\pi_{w_t}|\pi) \lesssim e^{-\frac{2c(t^\alpha-1)}{\alpha}} \frac{N}{2M} + \frac{c(N-S)(K+d)d}{4S(K-1)} \frac{1 + (1-\alpha) \log t}{t^{1-\alpha}}.$$

## C DETAILS OF EXPERIMENTS

### C.1 Model Specification

In this subsection, we describe the four examples (three real data and one synthetic) that we used for our experiments. For each model, we are given a set of points  $(x_n, y_n)_{n=1}^N$ , each consisting of features  $x_n \in \mathbb{R}^p$  and response  $y_n$ .

**Bayesian linear regression:** We use the model

$$\begin{aligned} [\beta \quad \log \sigma^2]^\top &\sim \mathcal{N}(0, I), \\ \forall n \in [N], y_n | x_n, \beta, \sigma^2 &\stackrel{\text{ind}}{\sim} \mathcal{N}([1 \quad x_n^\top] \beta, \sigma^2), \end{aligned}$$

where  $\beta \in \mathbb{R}^{p+1}$  is a vector of regression coefficients and  $\sigma^2 \in \mathbb{R}_+$  is the noise variance. Note that the prior here is not conjugate for the likelihood. The dataset<sup>2</sup> consists of flight delay information from  $N = 98,673$  observations. We study the difference, in minutes, between the scheduled and actual departure times against  $p = 10$  features including flight-specific and meteorological information.

**Bayesian logistic regression:** We use the model

$$\begin{aligned} \forall i \in [p+1], \quad \beta_i &\stackrel{\text{iid}}{\sim} \text{Cauchy}(0, 1), \\ \forall n \in [N], \quad y_n &\stackrel{\text{ind}}{\sim} \text{Bern}\left(\left(1 + \exp\left(-[1 \quad x_n^\top] \beta\right)\right)^{-1}\right), \end{aligned}$$

where  $\beta = [\beta_1 \quad \dots \quad \beta_{p+1}]^\top \in \mathbb{R}^{p+1}$  is a vector of regression coefficients. Here we use the same dataset as in linear regression, but instead model the relationship between whether a flight is cancelled using the same set of features. Note that of all flights included, only 0.058% were cancelled.

**Bayesian Poisson regression:** We use the model

$$\begin{aligned} \beta &\sim \mathcal{N}(0, I), \\ \forall n \in [N], y_n | x_n, \beta &\stackrel{\text{ind}}{\sim} \text{Pois}\left(\log\left(1 + e^{[1 \quad x_n^\top] \beta}\right)\right), \end{aligned}$$

where  $\beta \in \mathbb{R}^{p+1}$  is a vector of regression coefficients. Here we use a processed version of the bikeshare dataset<sup>3</sup> consisting of  $N = 15,641$  data points. We model the hourly count of rental bikes against  $p = 8$  features (e.g., temperature, humidity at the time, and whether or not the day is a workday).

**Bayesian sparse linear regression:** This is based on Example 4.1 from George and McCulloch (1993). We use the model

$$\begin{aligned} \sigma^2 &\sim \text{InvGam}(\nu/2, \nu\lambda/2), \\ \forall i \in [p], \quad \gamma_i &\stackrel{\text{iid}}{\sim} \text{Bern}(q), \\ \beta_i | \gamma_i &\stackrel{\text{ind}}{\sim} \mathcal{N}\left(0, (\mathbf{1}(\gamma_i = 0)\tau + \mathbf{1}(\gamma_i = 1)c\tau)^2\right), \\ \forall n \in [N], \quad y_n | x_n, \beta, \sigma^2 &\stackrel{\text{ind}}{\sim} \mathcal{N}(x_n^\top \beta, \sigma^2), \end{aligned}$$

where we set  $\nu = 0.1, \lambda = 1, q = 0.1, \tau = 0.1$ , and  $c = 10$ . Here we model the variance  $\sigma^2$ , the vector of regression coefficients  $\beta = [\beta_1 \quad \dots \quad \beta_p]^\top \in \mathbb{R}^p$  and the vector of binary variables  $\gamma = [\gamma_1 \quad \dots \quad \gamma_p]^\top \in \{0, 1\}^p$  indicating the inclusion of the  $p^{\text{th}}$  feature in the model. We set  $N = 50,000, p = 10, \beta^* = [0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 5 \quad 5 \quad 5 \quad 5 \quad 5]^\top$ , and generate a synthetic dataset by

$$\begin{aligned} \forall n \in [N], \quad x_n &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, I), \\ \epsilon_n &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, 25^2), \\ y_n &= x_n^\top \beta^* + \epsilon_n. \end{aligned}$$

For full-data inference results of the three real data examples, we ran Stan (Carpenter et al., 2017) with 10 parallel chains, each taking 15,000 steps where the first 5,000 were discarded, for a combined 100,000 draws. Full-data inference time was  $\sim 22$  minutes for linear regression,  $\sim 28$  minutes for logistic regression, and  $\sim 5$  minutes for Poisson regression, though note that these times are not comparable with the other results since Stan is highly optimized. For full-data inference result of the synthetic data example, we use the Gibbs sampler developed by George and McCulloch (1993) to generate 200,000 draws, with the first half discarded as burn-in. Full-data inference time for the sparse linear regression model was  $\sim 20$  minutes.

<sup>2</sup>The dataset is available at <https://github.com/NaitongChen/Sparse-Hamiltonian-Flows>.

<sup>3</sup>The dataset is available at <https://github.com/trevorcampbell/bayesian-coresets>.

## C.2 Parameter Settings

We begin by describing settings that apply across all experiments before moving to model-specific settings. For each method, we use 10,000 samples to estimate metrics that assess posterior approximation quality and sampling efficiency. For metrics that require information from the full-data posterior, we use a sample of size 100,000 from Stan (Carpenter et al., 2017) across all three real data experiments, and a sample of size 10,000 using the Gibbs sampler developed by George and McCulloch (1993) for the synthetic data experiment. To account for changes in  $w$ , for all coreset methods that use an MCMC kernel, we use the hit-and-run slice sampler with doubling (Bélisle et al., 1993; Neal, 2003) for linear and logistic regressions, the univariate slice sampler with doubling (Neal, 2003, Fig. 4) applied to each dimension for the more challenging Poisson regression, and the Gibbs sampler developed by George and McCulloch (1993) for sparse regression. Note that to get the 10,000 samples for metric evaluation from each method, we simulate 20,000 MCMC states and take the second half of the chain to ensure proper mixing. The exceptions are `CoresetMCMC` and `SHF`. For `CoresetMCMC`, no burn-in is necessary as in each of our experiments we adapt the coreset weights for more than 10,000 steps. For `SHF`, we have access to i.i.d. samples from the final trained posterior approximation.

Across all experiments, we set  $K = 2$  for `CoresetMCMC`. We follow Chen et al. (2022) and use 8 quasi-refreshment and 10 leapfrog steps between quasi-refreshments in `SHF`. To estimate the ELBO objective, we use a fresh minibatch of size 100 at each optimization iteration. Both `CoresetMCMC` and `SHF` are trained using ADAM (Kingma and Ba, 2014). For `QNC`, we use a sample of size 1,000 to estimate each weight update. We set the threshold to 0.05 for both `Austerity` and `Confidence`. For tuning `SGLD-CV` and `SGHMC-CV`, we follow Baker et al. (2019), and find the control variate using ADAM.

**Bayesian linear regression:** We train `CoresetMCMC` for 25,000 iterations with the ADAM step sizes set to 20, 1, 10, 10, 1, 1 for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ . We also train `CoresetMCMC-S` for 25,000 iterations with the ADAM step sizes set to  $10t^{-0.1}, 10t^{-0.1}, 10t^{-0.1}, 10t^{-0.3}, 20t^{-0.5}, 20t^{-0.3}$  for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ , where  $t$  is the iteration number. At each gradient estimate, we set the minibatch size to be 5 times the coreset size. We train `QNC` for 50 iterations so that the total number of samples drawn from the coreset posterior during optimization is the same between `CoresetMCMC` and `QNC`. In `QNC`, we choose the step size using a line search for each of the first 10 iterations. For the line search procedure, we use the curvature part of the Wolfe condition (Wolfe, 1969) to shrink the step size for a maximum of 20 times. For all iterations after, we use the same step size that is used in the  $10^{\text{th}}$  iteration. For `SHF`, we set the Adam step size to 0.002 and train the flow for 50,000 iterations. We use standard normal as the initial distribution across all dimensions.

For `Austerity`, we set the minibatch size to be 100. For both `Austerity` and `Confidence`, we use independent Gaussian proposal centred at the previous state with the variance of each component set to 0.01. For both `SGLD-CV` and `SGHMC-CV`, we set the subsample size to 500. For `SGLD-CV`, we use a constant step size of  $2 \times 10^{-5}$ . For `SGHMC-CV`, we use 30 leapfrog steps and a constant step size of  $8 \times 10^{-4}$ . To find the mode for the control variate, we use ADAM with default settings, drawing a subsamples of 100 observations until the weight updates for all dimensions are of magnitude smaller than  $5 \times 10^{-5}$ . All other parameters for the two stochastic gradient MCMC methods follow that of Baker et al. (2019).

**Bayesian logistic regression:** We train `CoresetMCMC` for 25,000 iterations with the ADAM step sizes set to 0.1, 5, 1, 1, 1, 0.1 for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ . We also train `CoresetMCMC-S` for 25,000 iterations with the ADAM step sizes set to  $10t^{-0.1}, 10t^{-0.1}, 10t^{-0.1}, 10t^{-0.3}, t^{-0.3}, 10t^{-0.5}$  for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ , where  $t$  is the iteration number. For both `Austerity` and `Confidence`, we use independent Gaussian proposal centred at the previous state with the variance of each component set to 0.001. To find the mode for the control variate for both stochastic gradient MCMC methods, we use ADAM with default settings, drawing a subsamples of 100 observations until the weight updates for all dimensions are of magnitude smaller than  $1 \times 10^{-3}$ . All other parameters are the same as in Bayesian linear regression.

To account for the class imbalance problem, for all coreset methods, we include all observations from the rare positive class if the coreset size is more than twice as big as the total number of observations with positive labels. Otherwise we sample our coreset to have 50% positive labels and 50% negative labels. This is only done for coreset methods because unlike subsampling MCMC methods which gets a fresh subsample each time, coreset points are only picked once.

**Bayesian Poisson regression:** We train `CoresetMCMC` for 50,000 iterations with the ADAM step sizes set to 2, 0.5, 0.5, 0.1, 0.05, 0.01 for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ . We also train `CoresetMCMC-S` for 50,000



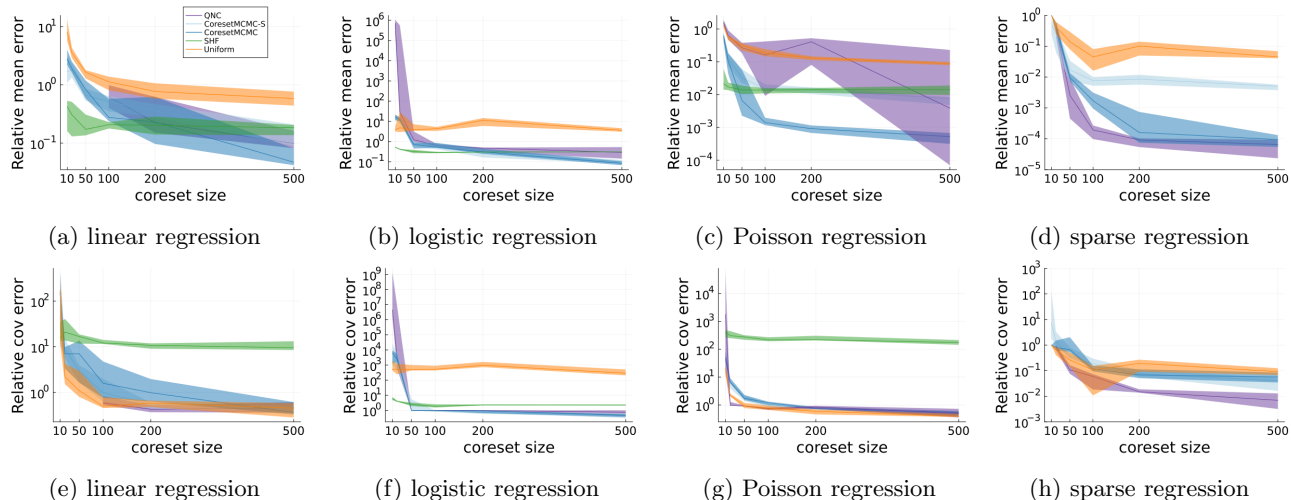


Figure 6: Comparison of coreset methods. Figs. 6a to 6d show posterior approximation quality via the relative mean error, and Figs. 6e to 6h show posterior approximation quality via the relative covariance error. The lines indicate the median, and error regions indicate 25<sup>th</sup> to 75<sup>th</sup> percentile from 10 runs.

iterations with the ADAM step sizes set to  $t^{-0.1}$ ,  $2t^{-0.1}$ ,  $2t^{-0.1}$ ,  $t^{-0.1}$ ,  $2t^{-0.3}$ ,  $0.5t^{-0.3}$  for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ , where  $t$  is the iteration number. At each gradient estimate, we set the minibatch size to be 10 times the coreset size. We train QNC for 100 iterations so that the total number of samples drawn from the coreset posterior during optimization is the same between CoresetMCMC and QNC. In QNC, we choose the step size using a line search for each of the first 10 iterations, where the step size may be shrunk up to 50 times each. For both *Austerity* and *Confidence*, we use independent Gaussian proposal centred at the previous state with the variance of each component set to 0.01. For SGLD-CV, we use a constant step size of  $2 \times 10^{-4}$ . For SGHMC-CV, we use 30 leapfrog steps and a constant step size of  $8 \times 10^{-4}$ . To find the mode for the control variate, we use ADAM with default settings, drawing a subsamples of 1000 observations until the weight updates for all dimensions are of magnitude smaller than  $1 \times 10^{-4}$ . All other parameters are the same as in Bayesian linear regression.

**Bayesian sparse linear regression:** We train CoresetMCMC for 25,000 iterations with the ADAM step sizes set to 0.1, 0.1, 1, 1, 0.1, 0.01 for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ . We also train CoresetMCMC-S for 25,000 iterations with the ADAM step sizes set to  $5t^{-0.1}$ ,  $10t^{-0.1}$ ,  $10t^{-0.1}$ ,  $t^{-0.1}$ ,  $5t^{-0.5}$ ,  $2t^{-0.5}$  for coreset sizes  $M = 10, 20, 50, 100, 200, 500$ , where  $t$  is the iteration number. At each gradient estimate, we set the minibatch size to be 10 times the coreset size. For the continuous components of both *Austerity* and *Confidence*, we use independent truncated Gaussian proposal that are centred at the previous state with the variance of each component set to 0.001. For the discrete components, we use independent Bernoulli proposals with the success probability set to 0.5. All other parameters are the same as in Bayesian linear regression. Note that due to the inclusion of discrete variables, SHF, SGLD-CV, and SGHMC-CV are no longer directly applicable, and are hence excluded from this experiment.

## D ADDITIONAL RESULTS

In this section we present additional experimental results that show the posterior relative mean error ( $\|\mu - \hat{\mu}\|_2 / \|\mu\|_2$ ) and relative covariance error ( $\|\Sigma - \hat{\Sigma}\|_F / \|\Sigma\|_F$ ) across all methods. Similar to the main text, here we use  $\hat{\mu}, \hat{\Sigma}$  to denote the mean and covariance estimated using draws from each method, and  $\mu, \Sigma$  are the same estimated for the full data posterior. Recall that the two-moment KL metric that we show in the main text combines both the posterior mean and covariance error into a single number.

Fig. 6 shows the relative mean and covariance error between coreset methods for all four models. Note that for the sparse regression model, this error is computed only on the continuous components. We see that the trends are mostly similar compared to those of two-moment KLs. However, we note that while SHF is able to capture the mean, it approximates the posterior covariances poorly, which is due to the limited expressiveness of the variational family. All other coreset methods involve using MCMC to sample from the coreset posterior, and so

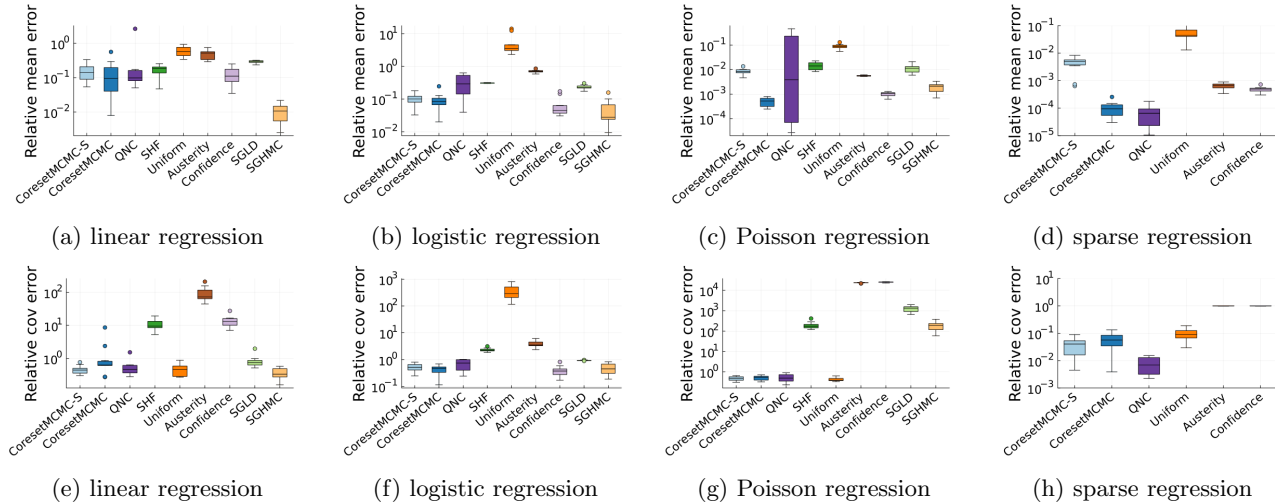


Figure 7: Comparison of coresets and subsampling MCMC methods. Figs. 7a to 7d show posterior approximation quality via the relative mean error, and Figs. 7e to 7h show posterior approximation quality via the relative covariance error. The boxplots indicate the median, 25<sup>th</sup>, and 75<sup>th</sup> percentiles from 10 runs.

the covariance estimates become closer to the true posterior as we increase the coresets size.

We now look at Fig. 7 to compare the relative mean and covariance error between coresets methods and subsampling MCMC methods. We set the coresets size  $M = 500$  for all coresets methods. We see that for models with only continuous variables, the lower two-moment KLs from SGHMC-CV that we show in the main text largely come from it being able to approximate the posterior mean well. However, in the Poisson regression example, stochastic optimization struggles to identify the mode of the posterior within the time for CoresetMCMC to train and sample. This causes the posterior approximation quality, especially that of the posterior covariance (as shown in Fig. 7g), to drop.